

EECS 127/227AT

Optimization Models in Engineering

Course Reader

Spring 2024

Acknowledgements

This reader is based on lectures from Spring 2021, Fall 2022, Spring 2023, and Fall 2023 iterations of EECS 127/227A by Prof. Gireeja Ranade. The reader was mostly written by Spring 2023 GSIs Druv Pai, Arwa Alanqary, and Aditya Ramabadran, and reviewed by Prof. Ranade. Fall 2022 tutor Jeffrey Wu collaborated with Druv on a writeup about the Eckart-Young theorem which was folded into the reader. Contributions from Prof. Venkat Anantharam and Chih-Yuan Chiu were also added in Fall 2023.

Contents

1	Introduction	4
1.1	What is Optimization?	4
1.2	Least Squares	7
1.3	Solution Concepts and Notation	10
1.4	(OPTIONAL) Infimum Versus Minimum	11
2	Linear Algebra Review	13
2.1	Norms	13
2.2	Gram-Schmidt and QR Decomposition	18
2.3	Fundamental Theorem of Linear Algebra	21
2.4	Symmetric Matrices	25
2.5	Principal Component Analysis	31
2.6	Singular Value Decomposition	34
2.7	Low-Rank Approximation	39
2.8	(OPTIONAL) Block Matrix Identities	46
3	Vector Calculus	48
3.1	Gradient, Jacobian, and Hessian	48
3.2	Taylor's Theorems	60
3.3	The Main Theorem	66
3.4	Directional Derivatives	67
3.5	(OPTIONAL) Matrix Calculus	68
4	Linear and Ridge Regression	70
4.1	Impact of Perturbations on Linear Regression	70
4.2	Ridge Regression	72
4.3	Principal Components Regression	73
4.4	Tikhonov Regression	75
4.5	Maximum Likelihood Estimation (MLE)	76
4.6	Maximum A Posteriori Estimation (MAP)	77
5	Convexity	79
5.1	Convex Sets	79
5.2	Convex Functions	101

5.3	Convex Optimization Problems	108
5.4	Solving Convex Optimization Problems	110
5.5	Problem Transformations and Reparameterizations	111
6	Gradient Descent	116
6.1	Strong Convexity and Smoothness	116
6.2	Gradient Descent	119
6.3	Variations: Stochastic Gradient Descent	125
6.4	Variations: Gradient Descent for Constrained Optimization	127
7	Duality	130
7.1	Lagrangian	130
7.2	Weak Duality	133
7.3	Strong Duality	137
7.4	Karush-Kuhn-Tucker (KKT) Conditions	141
7.5	(OPTIONAL) Conic Duality	144
8	Types of Optimization Problems	148
8.1	Linear Programs	148
8.2	Quadratic Programs	153
8.3	Quadratically-Constrained Quadratic Programs	156
8.4	Second-Order Cone Programs	156
8.5	Semidefinite Programming	164
8.6	General Taxonomy	171
9	Regularization and Sparsity	172
9.1	Recapping Ridge Regression and Defining LASSO	172
9.2	Understanding the Difference Between the ℓ^2 -Norm and the ℓ^1 -Norm	173
9.3	Analysis of LASSO Regression	176
9.4	Geometry of LASSO Regression	178
10	Advanced Descent Methods	180
10.1	Coordinate Descent	180
10.2	Newton's Method	183
10.3	Newton's Method with Linear Equality Constraints	185
10.4	(OPTIONAL) Interior Point Method	186
11	Applications	189
11.1	Deterministic Control and Linear-Quadratic Regulator	189
11.2	Support Vector Machines	193

Chapter 1

Introduction

Relevant sections of the textbooks:

- [1] Chapter 1.
- [2] Chapter 1.

1.1 What is Optimization?

Try to see what the following “problems” have in common.

- A statistical model, such as a neural network, trains using finite data samples.
- A robot learns a strategy using the environment, so that it does what you want.
- A major gas company decides what mixture of different fuels to process in order to get maximum profit.
- The EECS department decides how to set class sizes in order to maximize the number of credits offered subject to budget constraints.

While it might seem that these four examples are very distinct, they can all be formulated as minimizing an *objective function* over a *feasible set*. Thus, they can all be put into the framework of optimization.

To develop the basics of optimization, including precisely defining an objective function and a feasible set, we use some motivating examples from the third and fourth “problems”. (The first and second “problems” will be discussed at the very end of the course.)

Example 1 (Oil and Gas). Say that we are a gas company with 10^5 barrels of crude oil that we *must* refine by an expiration date. There are two refineries: one which processes crude oil into jet fuel, and one which processes crude oil into gasoline. We can sell a barrel of jet fuel to consumers for \$0.10, while we can sell a barrel of gasoline fuel for \$0.20. So, letting x_1 be a variable denoting the number of barrels of jet fuel produced, and x_2 be a variable denoting the number of barrels of gasoline produced, we aim to solve the problem:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{1}{10}x_1 + \frac{1}{5}x_2 \\ \text{s.t.} \quad & x_1 \geq 0 \\ & x_2 \geq 0 \end{aligned} \tag{1.1}$$

$$x_1 + x_2 = 10^5.$$

That is, we aim to choose x_1 and x_2 which maximize the *objective function* $\frac{1}{10}x_1 + \frac{1}{5}x_2$, but with the caveat that they must obey the *constraints* $x_1 \geq 0$, $x_2 \geq 0$, and $x_1 + x_2 = 10^5$. The *feasible set* is the set of all (x_1, x_2) pairs which obey the constraints. As you may have noticed, constraints can be equalities or inequalities in the x_i , which we formalize shortly.

The solution to this problem can be seen to be $(x_1^*, x_2^*) = (0, 10^5)$, which corresponds to refining all the crude oil into gasoline. This makes sense – after all, gasoline sells for more! And with all else equal between gasoline and jet fuel, to maximize our profit, we just need to produce gasoline.

To model another constraint, say that we need at least 10^3 gallons of jet fuel and $5 \cdot 10^2$ gallons of gasoline, we can directly incorporate them into the constraint set:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{1}{10}x_1 + \frac{1}{5}x_2 & (1.2) \\ \text{s.t.} \quad & x_1 \geq 0 \\ & x_2 \geq 0 \\ & x_1 \geq 10^3 \\ & x_2 \geq 5 \cdot 10^2 \\ & x_1 + x_2 = 10^5. \end{aligned}$$

We then notice that $x_1 \geq 0$ is made redundant by the constraint $x_1 \geq 10^3$. That is, no pair (x_1, x_2) which satisfies $x_1 \geq 10^3$ is not going to satisfy $x_1 \geq 0$. Thus, we can eliminate the latter constraint, since it defines the same feasible set. We can do the same thing for the constraints $x_2 \geq 0$ and $x_2 \geq 5 \cdot 10^2$, the latter making the former redundant. Thus, we can simplify the above problem to only include the redundant constraints:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{1}{10}x_1 + \frac{1}{5}x_2 & (1.3) \\ \text{s.t.} \quad & x_1 \geq 10^3 \\ & x_2 \geq 5 \cdot 10^2 \\ & x_1 + x_2 = 10^5. \end{aligned}$$

Let's say that we want to incorporate one final business need. Before, we were modeling that the oil refinement is free, since we don't have an objective or constraint term which involves this cost. Now, let us say that we can transport a total of $2 \cdot 10^6$ "barrel-miles" – that is, the number of barrels times the number of miles we can transport is no greater than $2 \cdot 10^6$. Let us further say that the jet fuel refinery is 10 miles away from the crude oil storage, and the gasoline refinery is 30 miles away from the crude oil storage. We can incorporate this further constraint into the constraint set directly:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{1}{10}x_1 + \frac{1}{5}x_2 & (1.4) \\ \text{s.t.} \quad & x_1 \geq 10^3 \\ & x_2 \geq 5 \cdot 10^2 \\ & 10x_1 + 30x_2 \leq 2 \cdot 10^6 \\ & x_1 + x_2 = 10^5. \end{aligned}$$

This is a good first problem; we have a non-trivial objective function, non-trivial inequality and equality constraints, and even got to work with manipulating constraints (so as to remove redundant ones)!

This type of optimization problem is called a *linear program*. We will learn more about how to formulate and solve linear programs later in the course.

A more generic reformulation of the above optimization problem is the following “standard form”.

Definition 2 (Standard Form of Optimization Problem)

We say that an optimization problem is written in *standard form* if it is of the form

$$\begin{aligned} \min_{\vec{x} \in \mathbb{R}^n} \quad & f_0(\vec{x}) \\ \text{s.t.} \quad & f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}. \end{aligned} \tag{1.5}$$

Here:

- $\vec{x} \in \mathbb{R}^n$ is the optimization variable.
- f_1, \dots, f_m and h_1, \dots, h_p are functions $\mathbb{R}^n \rightarrow \mathbb{R}$.
- f_0 is the objective function.
- f_i are inequality constraint functions; the expression “ $f_i(\vec{x}) \leq 0$ ” is an inequality constraint.
- Similarly, h_j are equality constraint functions, and the expression “ $h_j(\vec{x}) = 0$ ” is an equality constraint.
- The feasible set, i.e., the set of all \vec{x} that satisfy all constraints, is

$$\Omega \doteq \left\{ \vec{x} \in \mathbb{R}^n \mid \begin{array}{l} f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\} \end{array} \right\}. \tag{1.6}$$

We can thus also write the problem (1.5) as

$$\min_{\vec{x} \in \Omega} f_0(\vec{x}). \tag{1.7}$$

- A *solution* to this optimization problem is any $\vec{x}^* \in \Omega$ which attains the minimum value of $f_0(\vec{x})$ across all $\vec{x} \in \Omega$. Correspondingly, \vec{x}^* is also called a *minimizer* of f_0 over Ω .

It’s perfectly fine if $m = 0$ (in which case there are no inequality constraints) and/or $p = 0$ (in which case there are no equality constraints). If there are no constraints, then $\Omega = \mathbb{R}^n$ and the problem is called *unconstrained*; otherwise it is called *constrained*.

Let us try another example now, which has vector-valued quantities.

Example 3. Consider the following table of EECS courses:

Class	Size	Credits	Resources per Student
127	x_1	c_1	r_1
126	x_2	c_2	r_2
182	x_3	c_3	r_3
189	x_4	c_4	r_4
162	x_5	c_5	r_5
188	x_6	c_6	r_6
\vdots	\vdots	\vdots	\vdots

Suppose there are n classes in total. Let $\vec{x} \doteq [x_1 \ x_2 \ \cdots \ x_n]^\top \in \mathbb{R}^n$ be the decision variable, and let $\vec{c} \doteq [c_1 \ c_2 \ \cdots \ c_n]^\top \in \mathbb{R}^n$ and $\vec{r} \doteq [r_1 \ r_2 \ \cdots \ r_n]^\top \in \mathbb{R}^n$ be constants. Then, in order to maximize the total number of credit hours subject to a total resource budget b , we set up the linear program

$$\begin{aligned} \max_{\vec{x} \in \mathbb{R}^n} \quad & \vec{c}^\top \vec{x} \\ \text{s.t.} \quad & \vec{r}^\top \vec{x} \leq b \\ & x_i \geq 0, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \tag{1.8}$$

As notation, instead of the last set of constraints $x_i \geq 0$, we can write the vector constraint $\vec{x} \geq \vec{0}$.

More generally, recall that if we have a vector equality constraint $\vec{h}(\vec{x}) = \vec{0}$, it can be viewed as short-hand for the several scalar equality constraints $h_1(\vec{x}) = 0, \dots, h_p(\vec{x}) = 0$. Correspondingly, we define the vector inequality constraint $\vec{f}(\vec{x}) \leq \vec{0}$ to be short-hand for the several scalar inequality constraints $f_1(\vec{x}) \leq 0, \dots, f_m(\vec{x}) \leq 0$.

1.2 Least Squares

We begin with one of the simplest optimization problems, that of least squares. We've probably seen this formulation before. Mathematically, we are given a data matrix $A \in \mathbb{R}^{m \times n}$ and a vector of outcomes $\vec{y} \in \mathbb{R}^m$, and attempt to find a parameter vector $\vec{x} \in \mathbb{R}^n$ which minimizes the residual $\|A\vec{x} - \vec{y}\|_2^2$. Here $\|\cdot\|_2$ is the standard Euclidean norm $\|\vec{z}\|_2 \doteq \sqrt{\vec{z}^\top \vec{z}} = \sqrt{\sum_{i=1}^n z_i^2}$; it is labeled with the 2 for a reason we will see later in the course.

More precisely, we attempt to solve the following optimization problem:

$$\min_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{y}\|_2^2. \tag{1.9}$$

Theorem 4 (Least Squares Solution)

Let $A \in \mathbb{R}^{m \times n}$ have full column rank, and let $\vec{y} \in \mathbb{R}^m$. Then the solution to (1.9), i.e., the solution to


$$\min_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{y}\|_2^2, \tag{1.9}$$

is given by


$$\vec{x}^* = (A^\top A)^{-1} A^\top \vec{y}. \tag{1.10}$$

Proof. The idea is to find $A\vec{x} \in \mathcal{R}(A)$ which is closest to \vec{y} . Here $\mathcal{R}(A)$ is the range, or column space, or column span, of A . In general, We have no guarantee that $\vec{y} \in \mathcal{R}(A)$, so there is not necessarily an \vec{x} such that $A\vec{x} = \vec{y}$. Instead, we are finding an approximate solution to the equation $A\vec{x} = \vec{y}$.


Recall that $\mathcal{R}(A)$ is a subspace, and that \vec{y} itself may not belong to $\mathcal{R}(A)$. Thus we can visualize the geometry of the problem as the following picture:




We can now solve this problem using ideas from geometry. We claim that the closest point to \vec{y} contained in $\mathcal{R}(A)$ is the orthogonal projection of \vec{y} onto $\mathcal{R}(A)$; call this point \vec{z} . Also, define $\vec{e} \doteq \vec{y} - \vec{z}$. This gives the following diagram.



From this diagram, we see that \vec{e} is orthogonal to any vector in $\mathcal{R}(A)$. But remember that we still have to prove that \vec{z} is the closest point to \vec{y} within $\mathcal{R}(A)$. To see this, consider any another point $\vec{u} \in \mathcal{R}(A)$ and define $\vec{v} \doteq \vec{y} - \vec{u}$. This gives the following diagram:



To complete our proof, we define $\vec{w} \doteq \vec{z} - \vec{u}$, noting that the angle $\vec{u} \rightarrow \vec{z} \rightarrow \vec{y}$ is a right angle; in other words, \vec{w} and \vec{e} are orthogonal. This gives the following picture.



By the Pythagorean theorem, we see that

$$\|\vec{y} - \vec{u}\|_2^2 = \|\vec{v}\|_2^2 \quad (1.11)$$

$$= \|\vec{w}\|_2^2 + \|\vec{e}\|_2^2 \quad (1.12)$$

$$= \underbrace{\|\vec{z} - \vec{u}\|_2^2}_{>0} + \|\vec{e}\|_2^2 \quad (1.13)$$

$$> \|\vec{e}\|_2^2 \quad (1.14)$$

$$= \|\vec{y} - \vec{z}\|_2^2. \quad (1.15)$$

Therefore, \vec{z} is the closest point to \vec{y} within $\mathcal{R}(A)$.

Now, we want to find $\vec{z} \in \mathcal{R}(A)$, i.e., the orthogonal projection of \vec{y} onto $\mathcal{R}(A)$, such that $\vec{e} = \vec{y} - \vec{z}$ is orthogonal to all vectors in $\mathcal{R}(A)$. By the definition of $\mathcal{R}(A)$, it's equivalent to find $\vec{x}^* \in \mathbb{R}^n$ such that $\vec{y} - A\vec{x}^*$ is orthogonal to all vectors in $\mathcal{R}(A)$. Since the columns of A form a spanning set for $\mathcal{R}(A)$, it's equivalent to find $\vec{x}^* \in \mathbb{R}^n$ such that $\vec{y} - A\vec{x}^*$ is orthogonal to all columns of A . This implies

$$\vec{0} = A^\top (\vec{y} - A\vec{x}^*) \quad (1.16)$$

$$= A^\top \vec{y} - A^\top A\vec{x}^* \quad (1.17)$$

$$\implies A^\top A\vec{x}^* = A^\top \vec{y} \quad (1.18)$$

$$\implies \vec{x}^* = (A^\top A)^{-1} A^\top \vec{y}. \quad (1.19)$$

Here $A^\top A$ is invertible because A has full column rank. □


We'll conclude with a statistical application of least squares to linear regression. Suppose we are given data $(x_1, y_1), \dots, (x_n, y_n)$, and want to fit an *affine* model $y = mx + b$ through these data points. This corresponds to approximately solving the system

$$\begin{aligned} mx_1 + b &= y_1 \\ mx_2 + b &= y_2 \\ &\vdots \\ mx_n + b &= y_n. \end{aligned} \quad (1.20)$$

Formulating it in terms of vectors and matrices, we have

$$\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_n & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}. \quad (1.21)$$

In the case where the data is noisy or inconsistent with the model, as in the below figure, the linear system will be overdetermined and have no solutions. Then, we find an approximate solution – a line of best fit – via least squares on the above system.



As a last note, solving least squares (and similar problems) is easy because it is a so-called *convex* problem. Convex problems are easy to solve because any local optimum is a global optimum, which allows us to use a variety of simple techniques to find global optima. It is generally much more difficult to solve non-convex problems, though we solve a few during this course.

We discuss much more about convexity and convex problems later in the course.

1.3 Solution Concepts and Notation

Sometimes we assign values to our optimization problems. For example in the framework of (1.5) we may write

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \\ \text{s.t. } & f_i(\vec{x}) \leq 0 \quad \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0 \quad \forall j \in \{1, \dots, p\}. \end{aligned} \quad (1.22)$$

On the other hand, in the framework of (1.7) and using the definition of Ω in (1.6), we may write¹.

$$p^* = \min_{\vec{x} \in \Omega} f_0(\vec{x}). \quad (1.23)$$

This means that $p^* \in \mathbb{R}$ is the minimum value of f_0 over all $\vec{x} \in \Omega$; formally,

$$p^* = \min_{\vec{x} \in \Omega} f_0(\vec{x}) \doteq \min\{f_0(\vec{x}) \mid \vec{x} \in \Omega\}. \quad (1.24)$$

¹For the case where the minimum does not exist, but the infimum is finite, please see Section 1.4

As an example, consider the two-element set $\Omega = \{0, 1\}$ and $f_0(x) = 3x^2 + 2$. Then $p^* = \min\{f(0), f(1)\} = \min\{2, 5\} = 2$. We emphasize that p^* is a *real number*, not a vector.

To extract the minimizers, i.e., the points $\vec{x} \in \Omega$ which minimize $f_0(\vec{x})$, we use the argmin notation, which gives us the set of arguments which minimize our objective function. Formally, we define:

$$\operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x}) \doteq \left\{ \vec{x} \in \Omega \mid f_0(\vec{x}) = \min_{\vec{u} \in \Omega} f_0(\vec{u}) \right\} \quad (1.25)$$

We can thus write the set of solutions to (1.5) as

$$\begin{aligned} \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) & \quad (1.26) \\ \text{s.t. } f_i(\vec{x}) \leq 0 & \quad \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) = 0 & \quad \forall j \in \{1, \dots, p\}. \end{aligned}$$

And, as just discussed, we can write the set of solutions to (1.7) as

$$\operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x}). \quad (1.27)$$

We emphasize that the argmin is a *set of vectors*, any of which are an optimal solution, i.e., a minimizer, of the optimization problem at hand. It is possible for the argmin to contain 0 vectors (in which case the minimum value is not realized and the problem has no global optima), any positive number of vectors, or an infinite number of vectors.

Let us consider the same example as before. In particular, consider the two-element set $\Omega = \{0, 1\}$ and $f_0(x) = 3x^2 + 2$. Then $\operatorname{argmin}_{x \in \Omega} f_0(x) = \{0\}$. But, in different scenarios, the argmin can have zero elements; for example, if $f_0(x) = 3x$, then $\operatorname{argmin}_{x \in \mathbb{R}} f_0(x) = \emptyset$. And it can have multiple elements; for example, if $f_0(x) = 3x^2(x - 1)^2$, then $\operatorname{argmin}_{x \in \mathbb{R}} f_0(x) = \{0, 1\}$. It can even have infinitely many elements; for example, if $f_0(x) = 0$, then $\operatorname{argmin}_{x \in \mathbb{R}} f_0(x) = \mathbb{R}$.

Though we must remember to keep in mind that technically argmin is a set, in the problems we study, it usually contains exactly one element. Thus, instead of writing, for example, $\vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x})$, we may also write $\vec{x}^* = \operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x})$. The former expression is technically more correct, but both usages are fine, if — and only if — the argmin in question contains exactly one element.

1.4 (OPTIONAL) Infimum Versus Minimum

There is one remaining issue with our formulation, which we can conceptually consider as a “corner case”. What happens if the minimum does not exist? This may seem like a very esoteric case, yet one can construct a straightforward example, such as the following. We know that the minimum of any set of numbers must be contained in the set. But what happens if we try to find the minimum of the open interval $(0, 1)$? For any $x \in (0, 1)$ which we claim to be our minimum, we see that $\frac{x}{2}$ is also contained in $(0, 1)$ and is smaller than x , which is a contradiction to our claim. Thus the set $(0, 1)$ has no minimum.

It seems like 0 is a useful notion of “minimum” for this set — that is, it’s the largest number which is \leq all numbers in the set, i.e., its “greatest lower bound” — but it isn’t contained in the set and thus cannot be the minimum. Fortunately, this notion of greatest lower bound of a set is formalized in real analysis as the concept of an “infimum”, denoted \inf . For our purposes, we can think of the infimum as a generalization of the minimum which takes care of these corner cases and always exists. When the minimum exists, it is always equal to the infimum.

Based on this discussion, we can write our optimization problems as

$$p^* = \inf_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (1.28)$$

$$\begin{aligned} \text{s.t. } f_i(\vec{x}) &\leq 0 \quad \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) &= 0 \quad \forall j \in \{1, \dots, p\}. \end{aligned}$$

and

$$p^* = \inf_{\vec{x} \in \Omega} f_0(\vec{x}). \tag{1.29}$$

However, the argmin retains the same definition. In fact, one can prove that if we replaced the min in the argmin definition (1.25) with inf, that this “new” argmin would be exactly equivalent in every case to the “old” argmin, which we use henceforth. The analogous quantity to infimum for maximization — that is, the appropriate generalization of max — is the supremum, denoted sup.

Interested readers are encouraged to consult a real analysis textbook such as [3] for a more comprehensive coverage.

Though we have gone over the technical details here, for the rest of the course we will omit them for simplicity, and stick to using min and max (meaning inf and sup when the minimum and maximum do not exist).

Chapter 2

Linear Algebra Review

Relevant sections of the textbooks:

- [1] Appendix A.
- [2] Chapters 2, 3, 4, 5.

2.1 Norms

2.1.1 Definitions

Definition 5 (Norm)

Let \mathcal{V} be a vector space over \mathbb{R} . A function $f: \mathcal{V} \rightarrow \mathbb{R}$ is a norm if:

- Positive definiteness: $f(\vec{x}) \geq 0$ for all $\vec{x} \in \mathcal{V}$, and $f(\vec{x}) = 0$ if and only if $\vec{x} = \vec{0}$.
- Positive homogeneity: $f(\alpha\vec{x}) = |\alpha| f(\vec{x})$ for all $\alpha \in \mathbb{R}$ and $\vec{x} \in \mathcal{V}$.
- Triangle inequality: $f(\vec{x} + \vec{y}) \leq f(\vec{x}) + f(\vec{y})$ for all $\vec{x}, \vec{y} \in \mathcal{V}$.

We can check that the familiar Euclidean norm $\|\cdot\|_2: \vec{x} \mapsto \sqrt{\sum_{i=1}^n x_i^2}$ satisfies these properties. A generalization of the Euclidean norm is the following very useful class of norms.

Definition 6 (ℓ^p Norms)

Let $1 \leq p < \infty$. The ℓ^p -norm on \mathbb{R}^n is given by

$$\|\vec{x}\|_p \doteq \left(\sum_{i=1}^n |x_i|^p \right)^{1/p}. \quad (2.1)$$

The ℓ^∞ -norm on \mathbb{R}^n is given by

$$\|\vec{x}\|_\infty \doteq \max_{i \in \{1, \dots, n\}} |x_i|. \quad (2.2)$$

Example 7 (Examples of ℓ^p Norms).

- (a) The Euclidean norm, given by $\|\vec{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$, is an ℓ^p -norm for $p = 2$. (This is why we gave the subscript 2 to the Euclidean norm previously).
- (b) The ℓ^1 -norm is given by $\|\vec{x}\|_1 = \sum_{i=1}^n |x_i|$.
- (c) The ℓ^∞ -norm, given by $\|\vec{x}\|_\infty = \max_{i \in \{1, \dots, n\}} |x_i|$, is the limit of the ℓ^p norms as $p \rightarrow \infty$:

$$\|\vec{x}\|_\infty = \lim_{p \rightarrow \infty} \|\vec{x}\|_p. \quad (2.3)$$

We do not prove this here; it is left as an exercise.

2.1.2 Inequalities

There are a variety of useful *inequalities* which are associated with the ℓ^p norms. Before we provide them, we will take a second to discuss the importance of inequalities for optimization.

A priori, it may not be clear why we need to care about inequalities; why does it matter whether one arrangement of variables is always greater or less than another arrangement? It turns out that such inequalities are very helpful for characterizing the minimum and maximum of a given set of things; we can obtain upper bounds and lower bounds for things using these inequalities. This is definitely very helpful for optimization.

With that out of the way, let us get to the first major inequality.

Theorem 8 (Cauchy-Schwarz Inequality)

For any $\vec{x}, \vec{y} \in \mathbb{R}^n$, we have

$$|\vec{x}^\top \vec{y}| \leq \|\vec{x}\|_2 \|\vec{y}\|_2. \quad (2.4)$$

Proof. Let θ be the angle between \vec{x} and \vec{y} . We write

$$|\vec{x}^\top \vec{y}| = \|\vec{x}\|_2 \|\vec{y}\|_2 \cos \theta \quad (2.5)$$

$$= \|\vec{x}\|_2 \|\vec{y}\|_2 |\cos \theta| \quad (2.6)$$

$$\leq \|\vec{x}\|_2 \|\vec{y}\|_2. \quad (2.7)$$

□

We can get this result for ℓ^2 norms. A natural next question is whether we can generalize it to ℓ^p norms for $p \neq 2$. It turns out that we can, as we demonstrate shortly.

Theorem 9 (Hölder's Inequality)

Let $1 \leq p, q \leq \infty$ such that $\frac{1}{p} + \frac{1}{q} = 1$.^a Then for any $\vec{x}, \vec{y} \in \mathbb{R}^n$, we have

$$|\vec{x}^\top \vec{y}| \leq \sum_{i=1}^n |x_i y_i| \leq \|\vec{x}\|_p \|\vec{y}\|_q. \quad (2.8)$$

^aSuch pairs (p, q) are called *Hölder conjugates*.

This inequality collapses to [Cauchy-Schwarz Inequality](#) when $p = q = 2$. The proof is out of scope for now since it uses convexity.


Example 10 (Dual Norms). Fix $\vec{y} \in \mathbb{R}^n$. Let us solve the problem:

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \vec{x}^\top \vec{y}. \quad (2.9)$$

It is initially difficult to see how to proceed, so let us simplify the problem to get back onto familiar territory. We start with $p = 2$, so that the problem becomes:

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2 \leq 1}} \vec{x}^\top \vec{y}. \quad (2.10)$$

For $n = 2$, the feasible set and \vec{y} together look like the following:



For any $\vec{x} \in \mathbb{R}^n$, and θ the angle between \vec{x} and \vec{y} , we have

$$\vec{x}^\top \vec{y} = \|\vec{x}\|_2 \|\vec{y}\|_2 \cos \theta. \quad (2.11)$$


This term is maximized when $\cos \theta = 1$, or equivalently $\theta = 0$. Thus \vec{x} and \vec{y} must point in the same direction, i.e., \vec{x} is a scalar multiple of \vec{y} . And since we want to maximize this dot product, we must choose \vec{x} to maximize $\|\vec{x}\|_2$ subject to the constraint $\|\vec{x}\|_2 \leq 1$. Thus, we choose an \vec{x} which has $\|\vec{x}\|_2 = 1$ and points in the same direction as \vec{y} . This gives $\vec{x}^* = \vec{y} / \|\vec{y}\|_2$. Thus,

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2 \leq 1}} \vec{x}^\top \vec{y} = (\vec{x}^*)^\top \vec{y} = \left(\frac{\vec{y}}{\|\vec{y}\|_2} \right)^\top \vec{y} = \frac{\vec{y}^\top \vec{y}}{\|\vec{y}\|_2} = \frac{\|\vec{y}\|_2^2}{\|\vec{y}\|_2} = \|\vec{y}\|_2. \quad (2.12)$$

Now let us try $p = \infty$. The problem becomes

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_\infty \leq 1}} \vec{x}^\top \vec{y}. \quad (2.13)$$

The feasible set and \vec{y} are given by the following diagram.



Motivated by this diagram, we see that the constraint $\|\vec{x}\|_\infty \leq 1$ is equivalent to the $2n$ constraints $-1 \leq x_i$ and $x_i \leq 1$. Also, writing out the objective function

$$\vec{x}^\top \vec{y} = \sum_{i=1}^n x_i y_i = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n, \quad (2.14)$$

we see that the problem is

$$\begin{aligned} \max_{\vec{x} \in \mathbb{R}^n} & (x_1 y_1 + x_2 y_2 + \cdots + x_n y_n) \\ \text{s.t.} & -1 \leq x_i \leq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (2.15)$$

This problem has an interesting structure that will be repeated several times in the problems we discuss in this class. Namely, the objective function is the sum of several terms, each of which involves only one x_i . And the constraints are able to be partitioned into some groups, where the constraints in each group constrain only one x_i . Thus, this problem is *separable* into n different scalar problems, such that the optimal solutions for each scalar problem form an optimal solution for the vector problem. Namely, the problems are

$$\max_{\substack{x_i \in \mathbb{R} \\ -1 \leq x_i \leq 1}} x_i y_i \quad (2.16)$$

We solve this much simpler problem by hand. If $y_i > 0$ then $x_i^* = 1$; on the other hand, if $y_i \leq 0$ then $x_i^* = -1$. To summarize, $x_i^* = \text{sgn}(y_i)$, so that $x_i^* y_i = |y_i|$.


Putting all the scalar problems together, we see that $\vec{x}^* = \text{sgn}(\vec{y})$, and the vector problem's optimal value is given by

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_\infty \leq 1}} \vec{x}^\top \vec{y} = (\vec{x}^*)^\top \vec{y} = \sum_{i=1}^n x_i^* y_i = \sum_{i=1}^n \text{sgn}(y_i) y_i = \sum_{i=1}^n |y_i| = \|\vec{y}\|_1. \quad (2.17)$$

As a final exercise, we consider $p = 1$, so that the problem becomes

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_1 \leq 1}} \vec{x}^\top \vec{y}. \quad (2.18)$$

For $n = 2$, the feasible set and \vec{y} together look like the following:



We now bound the objective as

$$\vec{x}^\top \vec{y} \leq |\vec{x}^\top \vec{y}| \quad (2.19)$$

$$= \left| \sum_{i=1}^n x_i y_i \right| \quad (2.20)$$

$$\leq \sum_{i=1}^n |x_i y_i| \quad \text{by triangle inequality} \quad (2.21)$$

$$= \sum_{i=1}^n |x_i| |y_i| \quad (2.22)$$

$$\leq \sum_{i=1}^n |x_i| \left(\max_{i \in \{1, \dots, n\}} |y_i| \right) \quad (2.23)$$

$$= \left(\max_{i \in \{1, \dots, n\}} |y_i| \right) \sum_{i=1}^n |x_i| \quad (2.24)$$

$$= \|\vec{y}\|_\infty \|\vec{x}\|_1 \quad (2.25)$$

$$\leq \|\vec{y}\|_\infty. \quad (2.26)$$

Thus we have

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_1 \leq 1}} \vec{x}^\top \vec{y} \leq \|\vec{y}\|_\infty. \quad (2.27)$$

This inequality is actually an equality. To show this, we need to show the reverse inequality

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_1 \leq 1}} \vec{x}^\top \vec{y} \underset{\text{need to show}}{\geq} \|\vec{y}\|_\infty. \quad (2.28)$$

And showing *this* inequality amounts to choosing, for our fixed \vec{y} , a \vec{x} such that $\|\vec{x}\|_1 \leq 1$ and $\vec{x}^\top \vec{y} \geq \|\vec{y}\|_\infty$. This is also called “showing the maximum is attained”. To do this, we can find an \vec{x} such that $\|\vec{x}\|_p \leq 1$ and all the inequalities in the chain are met with equality.

- First, the inequality in (2.21) is a triangle inequality with the absolute value, i.e., $|\sum_{i=1}^n x_i y_i| \leq \sum_{i=1}^n |x_i y_i|$. To make sure this is an equality, it’s enough to make sure that all terms $x_i y_i$ are the same sign or 0.
- Next, the inequality in (2.23) says that $\sum_{i=1}^n |x_i| |y_i| \leq \sum_{i=1}^n |x_i| (\max_{i \in \{1, \dots, n\}} |y_i|)$. The most obvious instance in which this inequality is met with equality is when $|y_i| = \max_{j \in \{1, \dots, n\}} |y_j|$ for all i . But we can’t choose \vec{y} , as it’s fixed, so we can’t be assured that this holds. An alternate way in which this holds is that $|x_i| = 0$ for all i for which $|y_i| \neq \max_{j \in \{1, \dots, n\}} |y_j|$, i.e., $i \notin \operatorname{argmax}_{j \in \{1, \dots, n\}} |y_j|$.
- Finally, the inequality in (2.26) says that $\|\vec{x}\|_1 \|\vec{y}\|_\infty \leq \|\vec{y}\|_\infty$; to meet this inequality with equality, it is sufficient to have $\|\vec{x}\|_1 = 1$.

To meet all three of these constraints, we can construct \vec{x}^* via the following process:

- For each $i \notin \operatorname{argmax}_{j \in \{1, \dots, n\}} |y_j|$, set $\tilde{x}_i = 0$, as per the second bullet point above.
- For each $i \in \operatorname{argmax}_{j \in \{1, \dots, n\}} |y_j|$, set $\tilde{x}_i = \operatorname{sgn}(y_i)$, as per the first bullet point above.
- To get the true solution vector \vec{x}^* , divide $\tilde{\vec{x}}$ by $\|\tilde{\vec{x}}\|_1$; that is, $\vec{x}^* = \tilde{\vec{x}} / \|\tilde{\vec{x}}\|_1$. This ensures that $\|\vec{x}^*\|_1 = 1$, as per the third bullet point above.

This \vec{x}^* “achieves the maximum”, showing that

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_1 \leq 1}} \vec{x}^\top \vec{y} = \|\vec{y}\|_\infty. \quad (2.29)$$

This notion where the ℓ^2 -norm constraint leads to the ℓ^2 -norm objective, the ℓ^∞ -norm constraint leads to the ℓ^1 -norm objective, and the ℓ^1 -norm constraint leads to the ℓ^∞ -norm objective, hints at a greater pattern. Indeed, one can show that for $1 \leq p, q \leq \infty$ such that $\frac{1}{p} + \frac{1}{q} = 1$, an ℓ^p -norm constraint leads to an ℓ^q -norm objective:

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \vec{x}^\top \vec{y} = \|\vec{y}\|_q. \quad (2.30)$$

As before, we can prove this equality by proving the two constituent inequalities:

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \vec{x}^\top \vec{y} \leq \|\vec{y}\|_q \quad \text{and} \quad \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \vec{x}^\top \vec{y} \geq \|\vec{y}\|_q. \quad (2.31)$$

The proof of the first inequality (\leq) follows from applying Hölder's inequality to the objective function:

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \vec{x}^\top \vec{y} \leq \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \|\vec{x}\|_p \|\vec{y}\|_q = \|\vec{y}\|_q \cdot \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_p \leq 1}} \|\vec{x}\|_p = \|\vec{y}\|_q. \quad (2.32)$$

The second inequality (\geq) can follow if, for our fixed choice of \vec{y} , we produce some \vec{x} such that $\|\vec{x}\|_p \leq 1$ and $\vec{x}^\top \vec{y} \geq \|\vec{y}\|_q$, i.e., “the maximum is attained”. This is more complicated to do, and we won't do it here.

The above equality (2.30) means that the norms $\|\cdot\|_p$ and $\|\cdot\|_q$ are so-called *dual* norms. We will explore aspects of duality later in the course, though frankly we are just scratching the surface.

These problems, which are short and easy to state, contain a couple of core ideas within their solutions, which are broadly generalizable to a lot of optimization problems. For your convenience, we discuss these explicitly below.


Problem Solving Strategy 11 (Separating Vector Problems into Scalar Problems). *When trying to simplify an optimization problem, try to see if you can simplify it into several independent scalar problems. Then solve each scalar problem — this is usually much easier than solving the whole vector problem at once. The optimal solutions to each scalar problem will then form the optimal solution to the whole vector problem.*

Problem Solving Strategy 12 (Proving Optimality in an Optimization Problem). *To solve an optimization problem, you can use inequalities to bound the objective function, and then try to show that this bound is tight by finding a feasible choice of optimization variable which makes all the inequalities into equalities.*

2.2 Gram-Schmidt and QR Decomposition

The Gram-Schmidt algorithm is a way to turn a linearly independent set $\{\vec{a}_1, \dots, \vec{a}_k\}$ of vectors into an *orthonormal* set $\{\vec{q}_1, \dots, \vec{q}_k\}$ which spans the same space. To reiterate, an orthonormal set is a set of vectors in which each vector has norm 1 and is orthogonal to all others in the basis.

Suppose for simplicity that $n = k = 2$, and that we have the following vectors.




We begin with \vec{a}_1 . We want to construct a vector \vec{q}_1 such that

- it's orthogonal to all the \vec{q}_i which came before it — which is none of them, so we don't have to worry; and
- it has unit norm, so $\|\vec{q}_1\|_2 = 1$.

To achieve this, the simplest choice is

$$\vec{q}_1 \doteq \frac{\vec{a}_1}{\|\vec{a}_1\|_2}. \quad (2.33)$$




Then we go to \vec{a}_2 . To find \vec{q}_2 which is orthogonal to all the \vec{q}_i before it — that is, \vec{q}_1 — we subtract off the orthogonal projection of \vec{a}_2 onto \vec{q}_1 from \vec{a}_2 . The orthogonal projection of \vec{a}_2 onto \vec{q}_1 is given by

$$\vec{p}_2 \doteq \vec{q}_1(\vec{q}_1^\top \vec{a}_2) \quad (2.34)$$

and so the projection residual is given by


$$\vec{s}_2 \doteq \vec{a}_2 - \vec{p}_2 = \vec{a}_2 - \vec{q}_1(\vec{q}_1^\top \vec{a}_2). \quad (2.35)$$

Note that these formulas *only* hold because \vec{q}_1 is normalized, i.e., has norm 1.



While \vec{s}_2 is orthogonal to \vec{q}_1 , because we want a \vec{q}_2 that is normalized, we normalize \vec{s}_2 to get \vec{q}_2 :

$$\vec{q}_2 \doteq \frac{\vec{s}_2}{\|\vec{s}_2\|_2}. \quad (2.36)$$



If we had a vector \vec{a}_3 (and weren't limited by drawing in 2D space), we would ensure that \vec{q}_3 were orthogonal to \vec{q}_1 and \vec{q}_2 , as well as normalized, in a similar way as before. First we would compute the projection

$$\vec{p}_3 \doteq \vec{q}_1(\vec{q}_1^\top \vec{a}_3) + \vec{q}_2(\vec{q}_2^\top \vec{a}_3). \quad (2.37)$$

and the residual

$$\vec{s}_3 \doteq \vec{a}_3 - \vec{p}_3 = \vec{a}_3 - \vec{q}_1(\vec{q}_1^\top \vec{a}_3) - \vec{q}_2(\vec{q}_2^\top \vec{a}_3). \quad (2.38)$$

These projection formulas only hold because $\{\vec{q}_1, \vec{q}_2\}$ is an orthonormal set. And then we could compute

$$\vec{q}_3 \doteq \frac{\vec{s}_3}{\|\vec{s}_3\|_2}. \quad (2.39)$$

And so on. The general algorithm goes similar.

Algorithm 1 Gram-Schmidt algorithm.

```

1: function GRAMSCHMIDTALGORITHM(linearly independent set  $\{\vec{a}_1, \dots, \vec{a}_k\}$ )
2:    $\vec{q}_1 \doteq \vec{a}_1 / \|\vec{a}_1\|_2$ .
3:   for  $i \in \{2, 3, \dots, k\}$  do
4:      $\vec{p}_i \doteq \sum_{j=1}^{i-1} \vec{q}_j(\vec{q}_j^\top \vec{a}_i)$ 
5:      $\vec{s}_i \doteq \vec{a}_i - \vec{p}_i$ 
6:      $\vec{q}_i \doteq \vec{s}_i / \|\vec{s}_i\|_2$ 
7:   end for
8:   return orthonormal set  $\{\vec{q}_1, \dots, \vec{q}_k\}$ 
9: end function

```

This algorithm has the following two properties, which you can formally prove as an exercise.

Proposition 13 (Gram-Schmidt Algorithm)

Algorithm 1 has the following properties:

1. For each $i \in \{1, \dots, k\}$, we have

$$\text{span}(\vec{a}_1, \dots, \vec{a}_i) = \text{span}(\vec{q}_1, \dots, \vec{q}_i). \quad (2.40)$$

In particular, $\{\vec{a}_1, \dots, \vec{a}_k\}$ spans the same subspace as $\{\vec{q}_1, \dots, \vec{q}_k\}$, as was stated in our original goal.

2. $\{\vec{q}_1, \dots, \vec{q}_k\}$ is an orthonormal set.

The Gram-Schmidt algorithm leads to something called the QR decomposition. Because, for each i , we have $\text{span}(\vec{a}_1, \dots, \vec{a}_i) = \text{span}(\vec{q}_1, \dots, \vec{q}_i)$, it means that we can write \vec{a}_i as a linear combination of the \vec{q}_j :

$$\vec{a}_i = r_{1i}\vec{q}_1 + r_{2i}\vec{q}_2 + \dots + r_{ii}\vec{q}_i = \sum_{j=1}^i r_{ji}\vec{q}_j \quad (2.41)$$

Putting all the k equations in a matrix form, we can write

$$\begin{bmatrix} \vec{a}_1 & \dots & \vec{a}_k \end{bmatrix} = \begin{bmatrix} \vec{q}_1 & \dots & \vec{q}_k \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & \dots & r_{1k} \\ 0 & r_{22} & \dots & r_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & r_{kk} \end{bmatrix}. \quad (2.42)$$

More generally, we can decompose every tall matrix with full column rank into a product of a tall matrix with orthonormal columns Q and an upper-triangular matrix R .

Theorem 14 (QR Decomposition)

Let $A \in \mathbb{R}^{n \times k}$ where $k \leq n$ (so A is tall). Suppose A has full column rank. Then there is a matrix $Q \in \mathbb{R}^{n \times k}$ with orthonormal columns, and a matrix $R \in \mathbb{R}^{k \times k}$ which is upper triangular, such that $A = QR$.

As a final note, there are various alterations to the QR decomposition that work for matrices which are wide and/or do not have full column rank. Those are out of scope, but the idea is the same.

The QR decomposition is also relevant in numerical linear algebra, where it can be used to solve tall linear systems $A\vec{x} = \vec{y}$ efficiently, especially if the underlying matrix A has special structure. All such connections are out of scope.

2.3 Fundamental Theorem of Linear Algebra

The fundamental theorem of linear algebra is a tool for understanding what happens to vectors and vector spaces under a linear transformation. Matrix multiplication transforms one vector space into another. This is helpful for allowing us to change our coordinate system, which tells us more about the problem.

Definition 15 (Direct Sum)

Let $U, V \subseteq \mathbb{R}^n$ be subspaces. We say that U and V *direct sum* to \mathbb{R}^n , denoted $U \oplus V = \mathbb{R}^n$, if and only if:

- Every vector $\vec{x} \in \mathbb{R}^n$ can be written as $\vec{x} = \vec{x}_1 + \vec{x}_2$, where $\vec{x}_1 \in U$ and $\vec{x}_2 \in V$.
- Furthermore, this decomposition is unique, in the sense that if $\vec{x} = \vec{x}_1 + \vec{x}_2 = \vec{y}_1 + \vec{y}_2$ are two instances of the above decomposition, then $\vec{x}_1 = \vec{y}_1$ and $\vec{x}_2 = \vec{y}_2$.

Theorem 16 (Fundamental Theorem of Linear Algebra)

Let $A \in \mathbb{R}^{m \times n}$. Then

$$\mathcal{N}(A) \oplus \mathcal{R}(A^\top) = \mathbb{R}^n. \quad (2.43)$$

Note that we *cannot* replace $\mathcal{R}(A^\top)$ by $\mathcal{R}(A)$, since vectors in $\mathcal{R}(A)$ and $\mathcal{N}(A)$ do not even have the same number of entries or lie in the same Euclidean space. If we want to make a statement about $\mathcal{R}(A)$, we can replace A by A^\top in the above theorem to get the following corollary.

Corollary 17. Let $A \in \mathbb{R}^{m \times n}$. Then

$$\mathcal{N}(A^\top) \oplus \mathcal{R}(A) = \mathbb{R}^m. \quad (2.44)$$

To prove the fundamental theorem of linear algebra, we use a tool called the *orthogonal decomposition theorem*.

Definition 18 (Orthogonal Complement)

Let $S \subseteq \mathbb{R}^n$ be a subspace. The *orthogonal complement* of S , denoted S^\perp , is

$$S^\perp \doteq \{\vec{x} \in \mathbb{R}^n \mid \vec{s}^\top \vec{x} = 0 \text{ for all } \vec{s} \in S\} \quad (2.45)$$

Theorem 19 (Orthogonal Decomposition Theorem, Theorem 2.1 of [2])

Let $S \subseteq \mathbb{R}^n$ be a subspace. Then

$$S \oplus S^\perp = \mathbb{R}^n. \quad (2.46)$$

Proof. To prove this, we first need to prove the following claim:

Let $U, V \subseteq \mathbb{R}^n$ be subspaces. Then $U \oplus V = \mathbb{R}^n$ if and only if every vector $\vec{x} \in \mathbb{R}^n$ can be written as $\vec{x} = \vec{x}_1 + \vec{x}_2$, where $\vec{x}_1 \in U$ and $\vec{x}_2 \in V$, and $U \cap V = \{\vec{0}\}$.

To prove this claim, suppose first that $U \oplus V = \mathbb{R}^n$. Then every every vector $\vec{x} \in \mathbb{R}^n$ can be written as $\vec{x} = \vec{x}_1 + \vec{x}_2$, where $\vec{x}_1 \in U$ and $\vec{x}_2 \in V$. It remains to prove that $U \cap V = \{\vec{0}\}$. Suppose for the sake of contradiction that there exists $\vec{y} \neq \vec{0}$ such that $\vec{y} \in U \cap V$. Then

$$\vec{x} = (\vec{x}_1 + \vec{y}) + (\vec{x}_2 - \vec{y}). \quad (2.47)$$

Since $\vec{y} \in U$, we have $\vec{x}_1 + \vec{y} \in U$; since $\vec{y} \in V$, we have $\vec{x}_2 - \vec{y} \in V$. Thus

$$\vec{x} = \vec{x}_1 + \vec{x}_2 = (\vec{x}_1 + \vec{y}) + (\vec{x}_2 - \vec{y}) \quad (2.48)$$

are two distinct ways to write \vec{x} as the sum of vectors from U and V , so it cannot be true that $U \oplus V = \mathbb{R}^n$, a contradiction.

Towards the other direction, suppose that every vector $\vec{x} \in \mathbb{R}^n$ can be written as $\vec{x} = \vec{x}_1 + \vec{x}_2$, where $\vec{x}_1 \in U$ and $\vec{x}_2 \in V$, and $U \cap V = \{\vec{0}\}$. The only thing remaining to prove is that if

$$\vec{x} = \vec{x}_1 + \vec{x}_2 = \vec{z}_1 + \vec{z}_2 \quad (2.49)$$

where $\vec{x}_1, \vec{z}_1 \in U$ and $\vec{x}_2, \vec{z}_2 \in V$, then we must have $\vec{x}_1 = \vec{z}_1$ and $\vec{x}_2 = \vec{z}_2$. Suppose again for the sake of contradiction that there exists $\vec{x} \in \mathbb{R}^n$, $\vec{x}_1, \vec{z}_1 \in U$, and $\vec{x}_2, \vec{z}_2 \in V$ such that

$$\vec{x} = \vec{x}_1 + \vec{x}_2 = \vec{z}_1 + \vec{z}_2 \quad (2.50)$$

but $\vec{x}_1 \neq \vec{z}_1$ or $\vec{x}_2 \neq \vec{z}_2$. Then we have

$$\vec{0} = \vec{x} - \vec{x} = \vec{x}_1 + \vec{x}_2 - \vec{z}_1 - \vec{z}_2 = (\vec{x}_1 - \vec{z}_1) + (\vec{x}_2 - \vec{z}_2). \quad (2.51)$$

Thus, we have that

$$\vec{x}_1 - \vec{z}_1 = \vec{z}_2 - \vec{x}_2 \neq \vec{0}. \quad (2.52)$$

Since $\vec{x}_1, \vec{z}_1 \in U$, we have $\vec{x}_1 - \vec{z}_1 \in U$, and since $\vec{x}_2, \vec{z}_2 \in V$, we have $\vec{z}_2 - \vec{x}_2 \in V$. Since they are equal, we have $\vec{x}_1 - \vec{z}_1 \in U \cap V$ and nonzero. Thus $U \cap V \neq \{\vec{0}\}$, a contradiction.

This proves the above claim. Now to prove the actual theorem, we note that every vector $\vec{x} \in \mathbb{R}^n$ can be written as

$$\vec{x} = \text{proj}_S(\vec{x}) + (\vec{x} - \text{proj}_S(\vec{x})). \quad (2.53)$$

By definition, $\text{proj}_S(\vec{x}) \in S$, and because the projection residual is orthogonal to the subspace, we have $\vec{x} - \text{proj}_S(\vec{x}) \in S^\perp$. Thus every vector in \mathbb{R}^n can be written as the sum of a vector in S and S^\perp . It is an exercise to show that $S \cap S^\perp = \{\vec{0}\}$. Invoking the quoted claim completes the proof. \square

Using this theorem, the only thing we need to show to prove the fundamental theorem of linear algebra is that $\mathcal{N}(A)$ and $\mathcal{R}(A^\top)$ are orthogonal complements. We do this below.

Proof of Theorem 16. By Theorem 19, the only thing we need to show is that $\mathcal{N}(A) = \mathcal{R}(A^\top)^\perp$. This is a set equality; we show it by showing that $\mathcal{N}(A) \subseteq \mathcal{R}(A^\top)^\perp$ and that $\mathcal{N}(A) \supseteq \mathcal{R}(A^\top)^\perp$.

We first want to show that $\mathcal{N}(A) \subseteq \mathcal{R}(A^\top)^\perp$. That is, we want to show that for any $\vec{x} \in \mathcal{N}(A)$ we have $\vec{x} \in \mathcal{R}(A^\top)^\perp$. That is, for any $\vec{y} \in \mathcal{R}(A^\top)$, we want to show that $\vec{y}^\top \vec{x} = 0$.

Since $\vec{y} \in \mathcal{R}(A^\top)$ we can write $\vec{y} = A^\top \vec{w}$ for some $\vec{w} \in \mathbb{R}^m$. Then, since $\vec{x} \in \mathcal{N}(A)$ we have $A\vec{x} = \vec{0}$, so

$$\vec{y}^\top \vec{x} = (A^\top \vec{w})^\top \vec{x} \tag{2.54}$$

$$= \vec{w}^\top A\vec{x} \tag{2.55}$$

$$= \vec{w}^\top \vec{0} \tag{2.56}$$

$$= 0. \tag{2.57}$$

Thus \vec{x} and \vec{y} are orthogonal, so $\vec{x} \in \mathcal{R}(A^\top)^\perp$, which shows that $\mathcal{N}(A) \subseteq \mathcal{R}(A^\top)^\perp$.

We now want to show that $\mathcal{R}(A^\top)^\perp \subseteq \mathcal{N}(A)$. That is, we want to show that for any $\vec{x} \in \mathcal{R}(A^\top)^\perp$, we want to show that $\vec{x} \in \mathcal{N}(A)$. That is, we want to show that $A\vec{x} = \vec{0}$.

By definition, for every $\vec{y} \in \mathcal{R}(A^\top)$, we have $\vec{y}^\top \vec{x} = 0$. By writing $\vec{y} = A^\top \vec{w}$ for arbitrary $\vec{w} \in \mathbb{R}^m$, we get that for every $\vec{w} \in \mathbb{R}^m$ we have $(A^\top \vec{w})^\top \vec{x} = 0$. But the left-hand side is $\vec{w}^\top A\vec{x}$, so we have that $\vec{w}^\top A\vec{x} = 0$ for every $\vec{w} \in \mathbb{R}^m$. This is true for all $\vec{w} \in \mathbb{R}^m$, so it is true for the specific choice of $\vec{w} = A\vec{x}$, which yields

$$0 = \vec{w}^\top A\vec{x} \tag{2.58}$$

$$= (A\vec{x})^\top A\vec{x} \tag{2.59}$$

$$= \|A\vec{x}\|_2^2 \tag{2.60}$$

$$\implies A\vec{x} = \vec{0}. \tag{2.61}$$

This implies that $\vec{x} \in \mathcal{N}(A)$ as desired, so $\mathcal{R}(A^\top)^\perp \subseteq \mathcal{N}(A)$.

Thus, we have shown that $\mathcal{N}(A) = \mathcal{R}(A^\top)^\perp$, and so by Theorem 19 we have $\mathcal{N}(A) \oplus \mathcal{R}(A^\top) = \mathbb{R}^n$. \square

This will help us solve a very important optimization problem, which is considered “dual” to least squares in some sense. Recall that least squares helps us find an approximate solution to the linear system $A\vec{x} = \vec{y}$, when A is a *tall* matrix with full column rank. In other words, the linear system is over-determined, there are many more equations than unknowns, and there are generally no exact solutions, so we pick the solution with minimum squared error.

What about when A is a *wide* matrix with full *row* rank? There are now more unknowns than equations, and infinitely many exact solutions. So how do we pick one solution in particular? It really depends on which engineering problem we are solving. One *common* solution is to pick the minimum-energy or minimum-norm problem, which is the solution to the optimization problem:

$$\begin{aligned} \min_{\vec{x} \in \mathbb{R}^n} \quad & \|\vec{x}\|_2^2 \\ \text{s.t.} \quad & A\vec{x} = \vec{y}. \end{aligned} \tag{2.62}$$

Note that this principle of choosing the smallest or simplest solution — the “Occam’s Razor” principle — is much more broadly generalized beyond the case of finding solutions to linear systems, and is used within control theory and machine learning. But we deal with just this linear system case for now.

Theorem 20 (Minimum-Norm Solution)

Let $A \in \mathbb{R}^{m \times n}$ have full row rank, and let $\vec{y} \in \mathbb{R}^m$. Then the solution to Equation (2.62), i.e., the solution to

$$\begin{aligned} \min_{\vec{x} \in \mathbb{R}^n} \quad & \|\vec{x}\|_2^2 \\ \text{s.t.} \quad & A\vec{x} = \vec{y}, \end{aligned} \tag{2.62}$$

is given by

$$\vec{x}^* = A^\top (AA^\top)^{-1} \vec{y}. \tag{2.63}$$

Proof. Observe that the constraint $A\vec{x} = \vec{y}$ under-specifies the \vec{x} — in particular, any component of \vec{x} in $\mathcal{N}(A)$ will not affect the constraint and only the objective. In this sense, it is “wasteful”, and we should intuitively remove it. This motivates using Theorem 16 to decompose \vec{x} into a component inside $\mathcal{N}(A)$ — which we want to remove — and a component inside $\mathcal{R}(A^\top)$ — which we will optimize over.

Indeed, write $\vec{x} = \vec{u} + \vec{v}$, where $\vec{u} \in \mathcal{N}(A)$ and $\vec{v} \in \mathcal{R}(A^\top)$. Thus, there exists $\vec{w} \in \mathbb{R}^m$ such that $\vec{v} = A^\top \vec{w}$. The constraint becomes

$$\vec{y} = A\vec{x} \tag{2.64}$$

$$= A(\vec{u} + \vec{v}) \tag{2.65}$$

$$= A\vec{u} + A\vec{v} \tag{2.66}$$

$$= \vec{0} + AA^\top \vec{w} \tag{2.67}$$

$$= AA^\top \vec{w}. \tag{2.68}$$

And the objective function becomes

$$\|\vec{x}\|_2^2 = \|\vec{u} + \vec{v}\|_2^2 \tag{2.69}$$

$$= \vec{u}^\top \vec{u} + 2\vec{u}^\top \vec{v} + \vec{v}^\top \vec{v} \tag{2.70}$$

$$= \|\vec{u}\|_2^2 + 2\vec{v}^\top \vec{u} + \|\vec{v}\|_2^2 \tag{2.71}$$

$$= \|\vec{u}\|_2^2 + 2(A^\top \vec{w})^\top \vec{u} + \|\vec{v}\|_2^2 \tag{2.72}$$

$$= \|\vec{u}\|_2^2 + 2\vec{w}^\top A\vec{u} + \|\vec{v}\|_2^2 \tag{2.73}$$

$$= \|\vec{u}\|_2^2 + 2\vec{w}^\top \vec{0} + \|\vec{v}\|_2^2 \tag{2.74}$$

$$= \|\vec{u}\|_2^2 + 2 \cdot 0 + \|\vec{v}\|_2^2 \tag{2.75}$$

$$= \|\vec{u}\|_2^2 + \|\vec{v}\|_2^2 \tag{2.76}$$

$$= \|\vec{u}\|_2^2 + \|A^\top \vec{w}\|_2^2 \tag{2.77}$$

Thus, the minimum-norm problem can be reformulated in terms of \vec{u} and \vec{w} :

$$\begin{aligned} \min_{\substack{\vec{u} \in \mathbb{R}^n \\ \vec{w} \in \mathbb{R}^m}} \quad & \|\vec{u}\|_2^2 + \|A^\top \vec{w}\|_2^2 \\ \text{s.t.} \quad & \vec{y} = AA^\top \vec{w} \\ & A\vec{u} = \vec{0}. \end{aligned} \tag{2.78}$$

Now, because A has full row rank, AA^\top is invertible, so the first constraint implies that $\vec{w}^* = (AA^\top)^{-1} \vec{y}$, so $\vec{v}^* = A^\top \vec{w}^* = A^\top (AA^\top)^{-1} \vec{y}$. And because we are trying to minimize the objective, which only involves \vec{u} through $\|\vec{u}\|_2^2$, the ideal solution is to set $\vec{u}^* = \vec{0}$, which also satisfies the second constraint and so is feasible. Thus $\vec{x}^* = \vec{v}^* = A^\top (AA^\top)^{-1} \vec{y}$ as desired. \square

2.4 Symmetric Matrices

Symmetric matrices are a sub-class of matrices which have many special properties, and in engineering applications one usually tries to work with symmetric matrices as much as possible.

Definition 21 (Symmetric Matrix)

Let $A \in \mathbb{R}^{n \times n}$ be a square matrix. We say that A is *symmetric* if $A = A^\top$. The set of all symmetric matrices is denoted \mathbb{S}^n .


Equivalently, $A_{ij} = A_{ji}$ for all i and j .

Example 22. The 2×2 matrix $\begin{bmatrix} a & b \\ b & c \end{bmatrix}$ is symmetric.

Example 23 (Covariance Matrices). Any matrix of the form $A = BB^\top$, such as the covariance matrices we will discuss in the next section, is a symmetric matrix, since

$$A^\top = (BB^\top)^\top = (B^\top)^\top(B)^\top = BB^\top = A. \quad (2.79)$$

Example 24 (Adjacency Matrix). Consider an *undirected* connected graph $G = (V, E)$, for example the following:



Its adjacency matrix A has coordinate $A_{ij} = 1$ if $(i, j) \in E$, and $A_{ij} = 0$ otherwise; in the above example, we have

$$A = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}. \quad (2.80)$$

Since the graph is undirected, $(i, j) \in E$ if and only if $(j, i) \in E$, so $A_{ij} = A_{ji}$, and so A is a symmetric matrix.

Why do we care about symmetric matrices? Symmetric matrices have two nice properties: real eigenvalues, and guaranteed diagonalizability.

In general, a (non-symmetric) matrix need not be diagonalizable. For example, the matrix $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ is not diagonalizable. How can we characterize the diagonalizability of a matrix, then?

First, we will need the following definitions.

Definition 25 (Multiplicities)

Let $A \in \mathbb{R}^{n \times n}$, and let λ be an eigenvalue of A .

- (a) The *algebraic multiplicity* μ of eigenvalue λ in A is the number of times λ is a root of the characteristic polynomial $p_A(x) \doteq \det(xI - A)$ of A , i.e., it is the power of $(x - \lambda)$ in the factorization of $p_A(x)$.

(b) The *geometric multiplicity* ϕ of eigenvalue λ in A is the dimension of the null space $\Phi \doteq \mathcal{N}(\lambda I - A)$.

Theorem 26 (Diagonalizability)

A square matrix $A \in \mathbb{R}^{n \times n}$ is diagonalizable if and only if every eigenvalue of A has equal algebraic and geometric multiplicities.

Example 27 (Multiplicities of Degenerate Matrix). We were earlier told that the matrix $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ is not diagonalizable. To check this, let us compute its eigenvalues, algebraic multiplicities, and geometric multiplicities.

First, its characteristic polynomial is

$$p_A(x) = \det(xI - A) \quad (2.81)$$

$$= \det \left(\begin{bmatrix} x-1 & -1 \\ 0 & x-1 \end{bmatrix} \right) \quad (2.82)$$

$$= (x-1)^2. \quad (2.83)$$

Thus, A has only one eigenvalue $\lambda = 1$. Since $(x-1)$ has power 2 in the factorization of p_A , the eigenvalue $\lambda = 1$ has algebraic multiplicity $\mu = 2$.

The corresponding null space is

$$\Phi = \mathcal{N}(\lambda I - A) \quad (2.84)$$

$$= \mathcal{N} \left(\begin{bmatrix} 1-1 & -1 \\ 0 & 1-1 \end{bmatrix} \right) \quad (2.85)$$

$$= \mathcal{N} \left(\begin{bmatrix} 0 & -1 \\ 0 & 0 \end{bmatrix} \right) \quad (2.86)$$

$$= \text{span} \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad (2.87)$$

which has dimension $\phi = 1$. Thus, for $\lambda = 1$, we have $\mu \neq \phi$ and the matrix is indeed not diagonalizable.

This allows us to formally state the spectral theorem.

Theorem 28 (Spectral Theorem)

Let $A \in \mathbb{S}^n$ have eigenvalues λ_i with algebraic multiplicities μ_i , eigenspaces $\Phi_i \doteq \mathcal{N}(\lambda_i I - A)$, and geometric multiplicities $\phi_i \doteq \dim(\Phi_i)$.

- All eigenvalues are real: $\lambda_i \in \mathbb{R}$ for each i .
- Eigenspaces corresponding to different eigenvalues are orthogonal: Φ_i and Φ_j are orthogonal subspaces, i.e., for every $\vec{p}_i \in \Phi_i$ and $\vec{p}_j \in \Phi_j$ we have $\vec{p}_i^\top \vec{p}_j = 0$.
- A is diagonalizable: $\mu_i = \phi_i$ for each i .
- A is *orthonormally diagonalizable*; there exists an *orthonormal* matrix $U \in \mathbb{R}^{n \times n}$ and *diagonal* matrix $\Lambda \in \mathbb{R}^{n \times n}$ such that $A = U\Lambda U^\top$.

Recall that orthonormal matrices are matrices whose columns are orthonormal, i.e., are pairwise orthogonal and unit-norm. Orthonormal matrices U have the nice property that $U^T U = I$, and if U is square, then $U^T = U^{-1}$.

Proof of Theorem 28. Part (a) might be left to homework; part (b) will definitely be left to homework; we prove parts (c) and (d) here. In particular, we assume that parts (a) and (b) are true, and attempt to prove (d). Note that (d) implies (c), as an orthonormal diagonalization is a type of diagonalization, and so the existence of an orthonormal diagonalization must require the algebraic and geometric multiplicities to be equal.

Our proof strategy is to use induction on n , the size of the matrix. The base case of our induction is 1×1 matrices, for which the diagonalization is trivial. Now consider the inductive step. Our hope is, given $A \in \mathbb{S}^n$ which has eigenvalue λ , to get a decomposition of the form

$$A = V \begin{bmatrix} \lambda & \vec{0}^T \\ \vec{0} & B \end{bmatrix} V^T \quad \text{or equivalently} \quad V^T A V = \begin{bmatrix} \lambda & \vec{0}^T \\ \vec{0} & B \end{bmatrix} \quad (2.88)$$

where $V \in \mathbb{R}^{n \times n}$ is orthonormal and $B \in \mathbb{S}^{n-1}$ is symmetric. If we can do that, then we can inductively diagonalize $B = W \Gamma W^T$, and finally use that to construct a diagonalization for $A = U \Lambda U^T$, where $U \in \mathbb{R}^{n \times n}$ is orthonormal and $\Lambda \doteq \begin{bmatrix} \lambda & \vec{0}^T \\ \vec{0} & \Gamma \end{bmatrix}$.

Let \vec{u} be a unit-norm eigenvector of A corresponding to eigenvalue λ . Remember that we want an orthonormal matrix $V \in \mathbb{R}^{n \times n}$ which “isolates” λ . This motivates using \vec{u} and a basis of the orthogonal complement of $\text{span}(\vec{u})$ to form V . To construct this matrix V , we run Gram-Schmidt on the columns of the matrix $\begin{bmatrix} \vec{u} & I \end{bmatrix} \in \mathbb{R}^{n \times (n+1)}$, throwing out the single vector which will have 0 projection residual (there must be exactly one such vector by a counting argument; to get n linearly independent vectors from a spanning set of $n + 1$ vectors, we need to remove exactly one vector), and obtaining the orthonormal matrix $V = \begin{bmatrix} \vec{u} & V_1 \end{bmatrix} \in \mathbb{R}^{n \times n}$ where $V_1 \in \mathbb{R}^{n \times (n-1)}$ is itself orthonormal. By construction, we have $V_1^T \vec{u} = \vec{0}$ and $\vec{u}^T V_1 = \vec{0}^T$. Thus,

$$V^T A V = \begin{bmatrix} \vec{u} & V_1 \end{bmatrix}^T A \begin{bmatrix} \vec{u} & V_1 \end{bmatrix} \quad (2.89)$$

$$= \begin{bmatrix} \vec{u}^T \\ V_1^T \end{bmatrix} A \begin{bmatrix} \vec{u} & V_1 \end{bmatrix} \quad (2.90)$$

$$= \begin{bmatrix} \vec{u}^T \\ V_1^T \end{bmatrix} \begin{bmatrix} A\vec{u} & AV_1 \end{bmatrix} \quad (2.91)$$

$$= \begin{bmatrix} \vec{u}^T \\ V_1^T \end{bmatrix} \begin{bmatrix} \lambda\vec{u} & AV_1 \end{bmatrix} \quad (2.92)$$

$$= \begin{bmatrix} \lambda\vec{u}^T \vec{u} & \vec{u}^T AV_1 \\ \lambda V_1^T \vec{u} & V_1^T AV_1 \end{bmatrix} \quad (2.93)$$

$$= \begin{bmatrix} \lambda \|\vec{u}\|_2^2 & (A^T \vec{u})^T V_1 \\ \vec{0} & V_1^T AV_1 \end{bmatrix} \quad (2.94)$$

$$= \begin{bmatrix} \lambda & (A\vec{u})^T V_1 \\ \vec{0} & V_1^T AV_1 \end{bmatrix} \quad (2.95)$$

$$= \begin{bmatrix} \lambda & \lambda\vec{u}^T V_1 \\ \vec{0} & V_1^T AV_1 \end{bmatrix} \quad (2.96)$$

$$= \begin{bmatrix} \lambda & \vec{0}^T \\ \vec{0} & V_1^T AV_1 \end{bmatrix} \quad (2.97)$$

$$= \begin{bmatrix} \lambda & \vec{0}^\top \\ \vec{0} & B \end{bmatrix}, \quad (2.98)$$

where $B \doteq V_1^\top AV_1$ in accordance with our proof outline. Now we need to check that B is symmetric; indeed, we have

$$B^\top = (V_1^\top AV_1)^\top \quad (2.99)$$

$$= (V_1)^\top A^\top (V_1^\top)^\top \quad (2.100)$$

$$= V_1^\top AV_1 \quad (2.101)$$

$$= B. \quad (2.102)$$

By induction, we can orthonormally diagonalize this matrix as $B = W\Gamma W^\top \in \mathbb{R}^{(n-1) \times (n-1)}$, where $W \in \mathbb{R}^{(n-1) \times (n-1)}$ is orthonormal and $\Gamma \in \mathbb{R}^{(n-1) \times (n-1)}$ is diagonal. Thus, by using $W^{-1} = W^\top$, we have

$$\Gamma = W^\top BW \quad (2.103)$$

$$= W^\top V_1^\top AV_1 W \quad (2.104)$$

$$= (V_1 W)^\top A (V_1 W). \quad (2.105)$$

We want an orthonormal matrix $U \in \mathbb{R}^{n \times n}$ such that $U^\top AU = \Lambda = \begin{bmatrix} \lambda & \vec{0}^\top \\ \vec{0} & \Gamma \end{bmatrix}$. Thus, the above calculation motivates

the choice $U = \begin{bmatrix} \vec{u} & V_1 W \end{bmatrix} \in \mathbb{R}^{n \times n}$. Thus

$$U^\top AU = \begin{bmatrix} \vec{u} & V_1 W \end{bmatrix}^\top A \begin{bmatrix} \vec{u} & V_1 W \end{bmatrix} \quad (2.106)$$

$$= \begin{bmatrix} \vec{u}^\top \\ W^\top V_1^\top \end{bmatrix} A \begin{bmatrix} \vec{u} & V_1 W \end{bmatrix} \quad (2.107)$$

$$= \begin{bmatrix} \vec{u}^\top \\ W^\top V_1^\top \end{bmatrix} \begin{bmatrix} A\vec{u} & AV_1 W \end{bmatrix} \quad (2.108)$$

$$= \begin{bmatrix} \vec{u}^\top \\ W^\top V_1^\top \end{bmatrix} \begin{bmatrix} \lambda\vec{u} & AV_1 W \end{bmatrix} \quad (2.109)$$

$$= \begin{bmatrix} \lambda\vec{u}^\top\vec{u} & \vec{u}^\top AV_1 W \\ \lambda W^\top V_1^\top \vec{u} & W^\top V_1^\top AV_1 W \end{bmatrix} \quad (2.110)$$

$$= \begin{bmatrix} \lambda \|\vec{u}\|_2^2 & (A^\top \vec{u})^\top V_1 W \\ \lambda W^\top \vec{0} & W^\top BW \end{bmatrix} \quad (2.111)$$

$$= \begin{bmatrix} \lambda & (A\vec{u})^\top V_1 W \\ \vec{0} & \Gamma \end{bmatrix} \quad (2.112)$$

$$= \begin{bmatrix} \lambda & \lambda\vec{u}^\top V_1 W \\ \vec{0} & \Gamma \end{bmatrix} \quad (2.113)$$

$$= \begin{bmatrix} \lambda & \lambda\vec{0}^\top W \\ \vec{0} & \Gamma \end{bmatrix} \quad (2.114)$$

$$= \begin{bmatrix} \lambda & \vec{0}^\top \\ \vec{0} & \Gamma \end{bmatrix} \quad (2.115)$$

$$= \Lambda, \quad (2.116)$$

as desired. Thus $A = U\Lambda U^\top$ is an orthonormal diagonalization of A . This proves (d), and hence (c). \square

One nice thing about diagonalization is that we can read off the eigenvalues and eigenvectors from the components of the diagonalization.

Proposition 29

Let $A \in \mathbb{S}^n$ have orthonormal diagonalization $A = U\Lambda U^\top$, where $U = [\vec{u}_1 \ \cdots \ \vec{u}_n] \in \mathbb{R}^{n \times n}$ is square

orthonormal, and $\Lambda = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \in \mathbb{R}^{n \times n}$ is diagonal. Then for each i , the pair (λ_i, \vec{u}_i) is an eigenvalue-eigenvector pair for A .

Proof. By using $U^\top = U^{-1}$, we have

$$A = U\Lambda U^\top \quad (2.117)$$

$$AU = U\Lambda \quad (2.118)$$

$$A \begin{bmatrix} \vec{u}_1 & \cdots & \vec{u}_n \end{bmatrix} = \begin{bmatrix} \vec{u}_1 & \cdots & \vec{u}_n \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \quad (2.119)$$

$$\begin{bmatrix} A\vec{u}_1 & \cdots & A\vec{u}_n \end{bmatrix} = \begin{bmatrix} \lambda_1\vec{u}_1 & \cdots & \lambda_n\vec{u}_n \end{bmatrix}. \quad (2.120)$$

Therefore, each (λ_i, \vec{u}_i) is an eigenvalue-eigenvector pair of A . \square

Using this, we can work with another nice property of the orthonormal diagonalization. Namely, we can read off bases for $\mathcal{N}(A)$ and $\mathcal{R}(A)$. That is, a basis for $\mathcal{N}(A)$ is the set of eigenvectors \vec{u}_i corresponding to the eigenvalues λ_i of A which are equal to 0. Since U is orthonormal, the remaining eigenvectors \vec{u}_i span the orthogonal complement to $\mathcal{N}(A)$. But by the fundamental theorem of linear algebra (Theorem 16), we have $\mathcal{N}(A)^\perp = \mathcal{R}(A^\top) = \mathcal{R}(A)$, so these eigenvectors form a basis for $\mathcal{R}(A)$. Soon, we'll discover the singular value decomposition, which allows for this kind of decomposition of a matrix into its range and null spaces, except for arbitrary matrices.

Before we get into those, we will first state and solve a quick optimization problem which yields the eigenvalues of a symmetric matrix. This optimization problem turns out to be quite useful for further study of optimization.

Theorem 30 (Variational Characterization of Eigenvalues)

Let $A \in \mathbb{S}^n$. Let $\lambda_{\min}\{A\}$ and $\lambda_{\max}\{A\}$ be the maximum and minimum eigenvalues of A (which is well-defined since by the spectral theorem, all eigenvalues of A are real). Then

$$\lambda_{\max}\{A\} = \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \frac{\vec{x}^\top A \vec{x}}{\vec{x}^\top \vec{x}} = \max_{\|\vec{x}\|_2=1} \vec{x}^\top A \vec{x} \quad (2.121)$$

$$\lambda_{\min}\{A\} = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \frac{\vec{x}^\top A \vec{x}}{\vec{x}^\top \vec{x}} = \min_{\|\vec{x}\|_2=1} \vec{x}^\top A \vec{x}. \quad (2.122)$$

The term $\frac{\vec{x}^\top A \vec{x}}{\vec{x}^\top \vec{x}}$ is called the *Rayleigh quotient* of A ; it is a function of $\vec{x} \in \mathbb{R}^n$.

Proof. Before we start trying to prove any equalities, let us try to simplify the crucial term $\vec{x}^\top A \vec{x}$; the intuition behind this is that it looks like the easiest term to use the orthonormal diagonalization on and achieve results.

Let $A = U\Lambda U^\top$ be an orthonormal diagonalization of A . We have

$$\vec{x}^\top A \vec{x} = \vec{x}^\top U \Lambda U^\top \vec{x} \quad (2.123)$$

$$= (U^\top \vec{x})^\top \Lambda (U^\top \vec{x}) \quad (2.124)$$

$$= \vec{y}^\top \Lambda \vec{y} \quad (2.125)$$

$$= \sum_{i=1}^n \lambda_i \{A\} y_i^2 \quad (2.126)$$

with the invertible change of variables $\vec{y} \doteq U^\top \vec{x} \iff \vec{x} = U\vec{y}$. Also we note that this change of variables preserves the norm, i.e.,

$$\|\vec{y}\|_2^2 = \|U^\top \vec{x}\|_2^2 = \vec{x}^\top U U^\top \vec{x} = \vec{x}^\top \vec{x} = \|\vec{x}\|_2^2. \quad (2.127)$$

We now turn to the first equality chain (with max). Immediately, we have

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \frac{\vec{x}^\top A \vec{x}}{\vec{x}^\top \vec{x}} = \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \frac{\vec{x}^\top A \vec{x}}{\|\vec{x}\|_2^2} \quad (2.128)$$

$$= \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \left(\frac{\vec{x}}{\|\vec{x}\|_2} \right)^\top A \left(\frac{\vec{x}}{\|\vec{x}\|_2} \right). \quad (2.129)$$

Now, because the norm of our optimization variable \vec{x} does not matter, in that it only affects the objective through its normalization $\vec{x}/\|\vec{x}\|_2$, it is equivalent to optimize over only unit-norm \vec{x} , so

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{x} \neq \vec{0}}} \frac{\vec{x}^\top A \vec{x}}{\vec{x}^\top \vec{x}} = \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x}. \quad (2.130)$$

With the invertible change of variables $\vec{y} = U^\top \vec{x}$ already discussed, we can write

$$\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x} = \max_{\substack{\vec{y} \in \mathbb{R}^n \\ \|\vec{y}\|_2=1}} \sum_{i=1}^n \lambda_i y_i^2 \quad (2.131)$$

$$\leq \lambda_{\max}\{A\} \cdot \max_{\substack{\vec{y} \in \mathbb{R}^n \\ \|\vec{y}\|_2=1}} \sum_{i=1}^n y_i^2 \quad (2.132)$$

$$= \lambda_{\max}\{A\} \cdot \max_{\substack{\vec{y} \in \mathbb{R}^n \\ \|\vec{y}\|_2=1}} \|\vec{y}\|_2^2 \quad (2.133)$$

$$= \lambda_{\max}\{A\} \cdot \max_{\substack{\vec{y} \in \mathbb{R}^n \\ \|\vec{y}\|_2=1}} 1 \quad (2.134)$$

$$= \lambda_{\max}\{A\}. \quad (2.135)$$

It is left to exhibit a \vec{y} which makes this inequality an equality; indeed, it is achieved when $y_i = 1$ for one i such that $\lambda_i\{A\} = \lambda_{\max}\{A\}$ and $y_i = 0$ otherwise. The achieving \vec{x} can be recovered by $\vec{x} = U\vec{y}$. Note that since this \vec{y} is a standard basis vector $\vec{y} = \vec{e}_i$ for some i such that $\lambda_i\{A\} = \lambda_{\max}\{A\}$, then $\vec{x} = U\vec{y} = U\vec{e}_i = \vec{u}_i$, i.e., the i^{th} column of U , is an eigenvector of A corresponding to the maximum eigenvalue of A .

The analysis for $\lambda_{\min}\{A\}$ goes exactly analogously. \square

This characterization motivates defining a new sub-class (or really several new sub-classes) of matrices.

Definition 31 (Positive Semidefinite and Positive Definite Matrices)

Let $A \in \mathbb{S}^n$. We say that A is *positive semidefinite* (PSD), denoted $A \in \mathbb{S}_+^n$, if $\vec{x}^\top A \vec{x} \geq 0$ for all \vec{x} . We say that A is *positive definite* (PD), denoted $A \in \mathbb{S}_{++}^n$, if $\vec{x}^\top A \vec{x} > 0$ for all nonzero \vec{x} .

There are also negative semidefinite (NSD) and negative definite (ND) symmetric matrices, defined analogously. There are also indefinite symmetric matrices, which are none of the above. It is clear to see that PD matrices are themselves PSD.

Proposition 32

We have $A \in \mathbb{S}_+^n$ if and only if each eigenvalue of A is non-negative. Also, $A \in \mathbb{S}_{++}^n$ if and only if each eigenvalue of A is positive.

Proof. If A is PSD, we have

$$\lambda_{\min}\{A\} = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x} \geq 0. \quad (2.136)$$

Now suppose that each eigenvalue of A is non-negative. Then

$$0 \leq \lambda_{\min}\{A\} = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x} \quad (2.137)$$

which implies that $\vec{x}^\top A \vec{x} \geq 0$ for all \vec{x} with unit norm, and by scaling we see that $\vec{x}^\top A \vec{x} \geq 0$ for all $\vec{x} \neq \vec{0}$, while the inequality certainly holds for $\vec{x} = \vec{0}$. Thus $\vec{x}^\top A \vec{x} \geq 0$ for all \vec{x} so $A \in \mathbb{S}_+^n$.

If A is PD, we have

$$\lambda_{\min}\{A\} = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x} > 0. \quad (2.138)$$

Now suppose that each eigenvalue of A is positive. Then

$$0 < \lambda_{\min}\{A\} = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A \vec{x} \quad (2.139)$$

which implies that $\vec{x}^\top A \vec{x} > 0$ for all \vec{x} with unit norm, and by scaling we see that $\vec{x}^\top A \vec{x} > 0$ for all $\vec{x} \neq \vec{0}$, so $A \in \mathbb{S}_{++}^n$. \square

The final construction we discuss is that of the *positive semidefinite square root*.

Proposition 33

Let $A \in \mathbb{S}_+^n$. Then there exists a unique symmetric PSD matrix $B \in \mathbb{S}_+^n$, usually denoted $B = A^{1/2}$, such that $A = B^2$.

Proof. Discussion or homework. Note that there are non-symmetric matrices B such that $A = B^2$, but there is a unique PSD B . \square

2.5 Principal Component Analysis

Principal components analysis is a way to recover the eponymous principal components of the data. These principal components are those that are most representative of the data structure. Formally, if we have data in \mathbb{R}^d , we want to find an underlying p -dimensional linear structure, where $p \ll d$.

This idea has many, many use cases. For example, in modern machine learning, most data has thousands or millions of dimensions. In order to visualize it properly, we need to reduce its dimension to a reasonable number, in order to get an idea about the underlying structure of the data.

Let us first lay out some notation and definitions. Suppose we have the data points $\vec{x}_1, \dots, \vec{x}_n \in \mathbb{R}^d$. We organize these into a *data matrix* X where data points form the *rows*:¹


$$X \doteq \begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} \in \mathbb{R}^{n \times d} \quad \text{so that} \quad X^\top = \begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix} \in \mathbb{R}^{d \times n}. \quad (2.140)$$

We define the *covariance matrix* $C \in \mathbb{R}^{d \times d}$ by

$$C \doteq \frac{1}{n} X^\top X = \frac{1}{n} \begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix} \begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} = \frac{1}{n} \sum_{i=1}^n \vec{x}_i \vec{x}_i^\top. \quad (2.141)$$


We see that C is symmetric since $X^\top X$ is symmetric, so really $C \in \mathbb{S}^d$.

Before we progress further, let us try to get some intuition for what we might want to be doing. Consider the case $d = 2$ and $p = 1$. Suppose we have the dataset given as below.



While this dataset is clearly fully two-dimensional, there is equally clearly some inherent 1-dimensional linear structure to the data. So when we want to look for an underlying low-dimensional structure, we're looking for something like this. Here, if we could find the direction \vec{w} as in below:

¹Different textbooks handle this differently. For instance, some textbooks define a data matrix as one where the data points form the columns. If you're unsure, work it out from first principles.



And, if we are in generic \mathbb{R}^d space, we want to find orthonormal vectors $\vec{w}_1, \dots, \vec{w}_p$ such that projection onto them uncovers the underlying data structure. The process of accurately characterizing these \vec{w}_i is what we will discuss in what follows.

We begin with a motivating example. Consider the [MNIST dataset](#) of handwritten digits. Each image is a 28-pixel by 28-pixel grid with each numerical entry in the grid denoting the greyscale value at that grid point. This can be represented by a 28×28 size matrix, or alternatively unrolled into a $28^2 = 784$ -dimensional vector. It is impossible to directly visualize 784-dimensional space, so we seek to find $\vec{w}_1, \dots, \vec{w}_8 \in \mathbb{R}^{784}$ such that the projection onto the \vec{w}_i preserve a lot of structure. Say that we take $\vec{w}_i = \vec{e}_i$, where \vec{e}_i is the i^{th} standard basis vector in \mathbb{R}^{784} . Then for most images, the projection onto the \vec{w}_i 's will be 0 or near-0. Thus, the projection of all of the data onto the \vec{w}_i preserves almost none of the structure and collapses all points in the dataset to just a few points in \mathbb{R}^8 . There is instead a much more principled way to choose the \vec{w}_i that will preserve most of the structure.

We now discuss how to choose the first principal component $\vec{w}_1 \in \mathbb{R}^d$. To preserve the structure of the underlying data as much as possible, we want the vectors \vec{x}_i projected onto the span of \vec{w}_1 to be as close as possible to the original vectors \vec{x}_i . We also want $\|\vec{w}_1\|_2 = 1$. Thus, the error of the projection across all data points is

$$\text{err}(\vec{w}_1) = \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i - \vec{w}_1(\vec{w}_1^\top \vec{x}_i)\|_2^2.$$

Expanding, we have

$$\text{err}(\vec{w}_1) = \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i - \vec{w}_1(\vec{w}_1^\top \vec{x}_i)\|_2^2 \quad (2.142)$$

$$= \frac{1}{n} \sum_{i=1}^n (\vec{x}_i - \vec{w}_1(\vec{w}_1^\top \vec{x}_i))^\top (\vec{x}_i - \vec{w}_1(\vec{w}_1^\top \vec{x}_i)) \quad (2.143)$$

$$= \frac{1}{n} \sum_{i=1}^n (\vec{x}_i^\top \vec{x}_i - \vec{x}_i^\top \vec{w}_1(\vec{w}_1^\top \vec{x}_i) - (\vec{w}_1^\top \vec{x}_i)^\top \vec{x}_i + (\vec{w}_1^\top \vec{x}_i)^\top (\vec{w}_1(\vec{w}_1^\top \vec{x}_i))) \quad (2.144)$$

$$= \frac{1}{n} \sum_{i=1}^n (\vec{x}_i^\top \vec{x}_i - 2(\vec{x}_i^\top \vec{w}_1)^2 + (\vec{w}_1^\top \vec{w}_1)(\vec{w}_1^\top \vec{x}_i)^2) \quad (2.145)$$

$$= \frac{1}{n} \sum_{i=1}^n (\|\vec{x}_i\|_2^2 - 2(\vec{x}_i^\top \vec{w}_1)^2 + (\vec{w}_1^\top \vec{x}_i)^2) \quad (2.146)$$

$$= \frac{1}{n} \sum_{i=1}^n (\|\vec{x}_i\|_2^2 - (\vec{x}_i^\top \vec{w}_1)^2). \quad (2.147)$$

Now solving the principal components optimization problem gives

$$\min_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \text{err}(\vec{w}_1) = \min_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \frac{1}{n} \sum_{i=1}^n (\|\vec{x}_i\|_2^2 - (\vec{x}_i^\top \vec{w}_1)^2) \quad (2.148)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 + \min_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \frac{1}{n} \sum_{i=1}^n (-(\vec{x}_i^\top \vec{w}_1)^2) \quad (2.149)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \max_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \frac{1}{n} \sum_{i=1}^n (\vec{x}_i^\top \vec{w}_1)^2 \quad (2.150)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \max_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \frac{1}{n} \sum_{i=1}^n \vec{w}_1^\top \vec{x}_i \vec{x}_i^\top \vec{w}_1 \quad (2.151)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \max_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \vec{w}_1^\top \left(\frac{1}{n} \sum_{i=1}^n \vec{x}_i \vec{x}_i^\top \right) \vec{w}_1 \quad (2.152)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \max_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \vec{w}_1^\top \left(\frac{1}{n} X^\top X \right) \vec{w}_1 \quad (2.153)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \max_{\substack{\vec{w}_1 \in \mathbb{R}^d \\ \|\vec{w}_1\|_2=1}} \vec{w}_1^\top C \vec{w}_1 \quad (2.154)$$

$$= \frac{1}{n} \sum_{i=1}^n \|\vec{x}_i\|_2^2 - \lambda_{\max}\{C\} \quad (2.155)$$

with the \vec{w}_1 achieving this upper bound being the eigenvector \vec{u}_{\max} corresponding to the eigenvalue $\lambda_{\max}\{C\}$. Thus, the first principal component is exactly an eigenvector corresponding to the largest eigenvalue of the dot product matrix $C = X^\top X/n$.

This computation is a special case of the singular value decomposition, which is used in practice to compute the PCA of a dataset; understanding this decomposition will allow us to neatly compute the other principal components (i.e., second, third, fourth,...), as well.

2.6 Singular Value Decomposition

Definition 34 (SVD)

Let $A \in \mathbb{R}^{m \times n}$ have rank r . A *singular value decomposition (SVD)* of A is a decomposition of the form

$$A = U \Sigma V^\top \quad (2.156)$$

$$= \begin{bmatrix} U_r & U_{m-r} \end{bmatrix} \begin{bmatrix} \Sigma_r & 0_{r \times (n-r)} \\ 0_{(m-r) \times r} & 0_{(m-r) \times (n-r)} \end{bmatrix} \begin{bmatrix} V_r^\top \\ V_{n-r}^\top \end{bmatrix} \quad (2.157)$$

$$= U_r \Sigma_r V_r^\top \quad (2.158)$$

$$= \sum_{i=1}^r \sigma_i \vec{u}_i \vec{v}_i^\top, \quad (2.159)$$

where:

- $U \in \mathbb{R}^{m \times m}$, $U_r \in \mathbb{R}^{m \times r}$, $U_{m-r} \in \mathbb{R}^{m \times (m-r)}$, $V \in \mathbb{R}^{n \times n}$, $V_r \in \mathbb{R}^{n \times r}$, and $V_{n-r} \in \mathbb{R}^{n \times (n-r)}$ are orthonormal matrices, where $U = \begin{bmatrix} U_r & U_{m-r} \end{bmatrix}$ has columns $\vec{u}_1, \dots, \vec{u}_m$ (*left singular vectors*) and $V = \begin{bmatrix} V_r & V_{n-r} \end{bmatrix}$ has columns $\vec{v}_1, \dots, \vec{v}_n$ (*right singular vectors*).

- $\Sigma_r = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \in \mathbb{R}^{r \times r}$ is a diagonal matrix with ordered positive entries $\sigma_1 \geq \dots \geq \sigma_r > 0$ (*singular values*), and the zero matrices in the $\Sigma = \begin{bmatrix} \Sigma_r & 0_{r \times (n-r)} \\ 0_{(m-r) \times r} & 0_{(m-r) \times (n-r)} \end{bmatrix}$ matrix are shaped to ensure that $\Sigma \in \mathbb{R}^{m \times n}$.

Suppose that A is tall (so $m > n$) with full column rank n . Then the SVD looks like the following:

$$A = U \begin{bmatrix} \Sigma_n \\ 0_{(m-n) \times n} \end{bmatrix} V^\top. \quad (2.160)$$

On the other hand, if A is wide (so $m < n$) with full row rank m , then the SVD looks like the following:

$$A = U \begin{bmatrix} \Sigma_m & 0_{m \times (n-m)} \end{bmatrix} V^\top. \quad (2.161)$$

The last (summation) form of the SVD is called the *dyadic SVD*; this is because terms of the form $\vec{p}\vec{q}^\top$ are called *dyads*, and the dyadic SVD expresses the matrix A as the sum of dyads.

All forms of the SVD are useful conceptually and computationally, depending on the problem we are working on.

We now discuss a method to construct the SVD. Suppose $A \in \mathbb{R}^{m \times n}$ has rank r . We consider the symmetric matrix $A^\top A$ which has rank r and thus r nonzero eigenvalues, which are positive. We can order its eigenvalues as $\lambda_1 \geq \dots \geq \lambda_r > \lambda_{r+1} = \dots = \lambda_n = 0$, say with corresponding orthonormal eigenvectors $\vec{v}_1, \dots, \vec{v}_n$.

Then, for $i \in \{1, \dots, r\}$, we define $\sigma_i \doteq \sqrt{\lambda_i} > 0$, and $\vec{u}_i \doteq A\vec{v}_i/\sigma_i$. This only gives us r vectors \vec{u}_i , but we need m of them to construct $U \in \mathbb{R}^{m \times m}$. To find the remaining \vec{u}_i we use Gram-Schmidt on the matrix $\begin{bmatrix} \vec{u}_1 & \dots & \vec{u}_r & I \end{bmatrix} \in \mathbb{R}^{m \times (r+m)}$, throwing out the r vectors whose projection residual onto previously processed vectors is 0.

More formally, we can write an algorithm:

Algorithm 2 Construction of the SVD.

function SVD($A \in \mathbb{R}^{m \times n}$)

$r \doteq \text{rank}(A)$

$(\lambda_1, \vec{v}_1), \dots, (\lambda_n, \vec{v}_n) \doteq \text{EIGENPAIRS}(A^\top A) \triangleright \lambda_1 \geq \dots \geq \lambda_r > \lambda_{r+1} = \dots = \lambda_n = 0$, and \vec{v}_i orthonormal.

for $i \in \{1, \dots, r\}$ **do**

$\sigma_i \doteq \sqrt{\lambda_i}$

$\triangleright \sigma_i > 0$

$\vec{u}_i \doteq A\vec{v}_i/\sigma_i$

end for

$\vec{u}_1, \dots, \vec{u}_r, \vec{u}_{r+1}, \dots, \vec{u}_m \doteq \text{GRAMSCHMIDT}(\begin{bmatrix} \vec{u}_1 & \dots & \vec{u}_r & I \end{bmatrix})$

return $\{\vec{u}_1, \dots, \vec{u}_m\}, \{\sigma_1, \dots, \sigma_r\}, \{\vec{v}_1, \dots, \vec{v}_n\}$

end function

It's clear that Algorithm 2 gives an orthonormal basis $\{\vec{v}_1, \dots, \vec{v}_n\}$ for \mathbb{R}^n that can be constructed into the orthonormal V matrix, and that it gives singular values $\sigma_1 \geq \dots \geq \sigma_r > 0$. We aim to show two things: the $\{\vec{u}_1, \dots, \vec{u}_m\}$ are orthonormal, and that $A = U\Sigma V^\top$ where U , Σ , and V are constructed using the returned vectors and scalars.

Proposition 35

In the context of Algorithm 2, $\{\vec{u}_1, \dots, \vec{u}_m\}$ is an orthonormal set.

Proof. From our invocation of Gram-Schmidt, $\{\vec{u}_{r+1}, \dots, \vec{u}_m\}$ is an orthonormal set which spans an orthogonal subspace to the span of $\{\vec{u}_1, \dots, \vec{u}_r\}$. Thus, we need to show that $\{\vec{u}_1, \dots, \vec{u}_r\}$ are orthonormal.

Indeed, take $1 \leq i < j \leq r$. Then since the \vec{v}_j are orthonormal eigenvectors of $A^\top A$, we have

$$\vec{u}_i^\top \vec{u}_j = \left(\frac{A\vec{v}_i}{\sigma_i} \right)^\top \left(\frac{A\vec{v}_j}{\sigma_j} \right) \quad (2.162)$$

$$= \frac{\vec{v}_i^\top A^\top A \vec{v}_j}{\sigma_i \sigma_j} \quad (2.163)$$

$$= \frac{\lambda_j \vec{v}_i^\top \vec{v}_j}{\sigma_i \sigma_j} \quad (2.164)$$

$$= \frac{\lambda_j}{\sigma_i \sigma_j} \underbrace{\vec{v}_i^\top \vec{v}_j}_{=0} \quad (2.165)$$

$$= 0. \quad (2.166)$$

On the other hand, for a specific $i \in \{1, \dots, r\}$, using that $\sigma_i^2 = \lambda_i$, we have

$$\|\vec{u}_i\|_2^2 = \left\| \frac{A\vec{v}_i}{\sigma_i} \right\|_2^2 \quad (2.167)$$

$$= \left(\frac{A\vec{v}_i}{\sigma_i} \right)^\top \left(\frac{A\vec{v}_i}{\sigma_i} \right) \quad (2.168)$$

$$= \frac{\vec{v}_i^\top A^\top A \vec{v}_i}{\sigma_i^2} \quad (2.169)$$

$$= \frac{\lambda_i \vec{v}_i^\top \vec{v}_i}{\sigma_i^2} \quad (2.170)$$

$$= \frac{\lambda_i}{\sigma_i^2} \underbrace{\vec{v}_i^\top \vec{v}_i}_{=1} \quad (2.171)$$

$$= 1. \quad (2.172)$$

Thus the set $\{\vec{u}_1, \dots, \vec{u}_r\}$ is orthonormal, so the whole set $\{\vec{u}_1, \dots, \vec{u}_m\}$ is orthonormal. \square

Proposition 36

In the context of Algorithm 2, we have $A = U\Sigma V^\top$.

Proof. By construction, we have

$$A\vec{v}_i = \sigma_i \vec{u}_i \quad \text{for all } i \in \{1, \dots, r\}, \quad (2.173)$$

$$A\vec{v}_i = \vec{0} \quad \text{for all } i \in \{r+1, \dots, m\}, \quad (2.174)$$

This gives us

$$AV = A \begin{bmatrix} \vec{v}_1 & \cdots & \vec{v}_r & \vec{v}_{r+1} & \cdots & \vec{v}_m \end{bmatrix} \quad (2.175)$$

$$= \begin{bmatrix} A\vec{v}_1 & \cdots & A\vec{v}_r & A\vec{v}_{r+1} & \cdots & A\vec{v}_n \end{bmatrix} \quad (2.176)$$

$$= \begin{bmatrix} \sigma_1 \vec{u}_1 & \cdots & \sigma_r \vec{u}_r & \vec{0} & \cdots & \vec{0} \end{bmatrix} \quad (2.177)$$

$$= \begin{bmatrix} U_r \Sigma_r & 0 \end{bmatrix} \quad (2.178)$$

On the other hand, we have

$$U\Sigma = \begin{bmatrix} U_r & U_{m-r} \end{bmatrix} \begin{bmatrix} \Sigma_r & 0 \\ 0 & 0 \end{bmatrix} \quad (2.179)$$

$$= \begin{bmatrix} U_r \Sigma_r + U_{m-r} \cdot 0 & U_r \cdot 0 + U_{m-r} \cdot 0 \end{bmatrix} \quad (2.180)$$


$$= \begin{bmatrix} U_r \Sigma_r & 0 \end{bmatrix}. \quad (2.181)$$

Thus $AV = U\Sigma$. Since V is orthonormal, $V^\top = V^{-1}$, so $A = U\Sigma V^\top$. \square

The SVD is not unique: the Gram-Schmidt process could have used any basis for \mathbb{R}^m that wasn't the columns of I and still have been valid; if you had multiple eigenvectors of $A^\top A$ with the same eigenvalue then the choice of eigenvectors in the diagonalization would not be unique; and even if you didn't have multiple eigenvectors with the same eigenvalue, the eigenvectors would only be determined up to a sign change $\vec{v} \mapsto -\vec{v}$ anyways. So there is a lot of ambiguity in the construction, which reflects the non-uniqueness.


We now discuss the geometry of the SVD, especially how each component of the SVD acts on vectors. To do this we will fix $A \in \mathbb{R}^{2 \times 2}$ with SVD $A = U\Sigma V^\top$, and find the behavior of $A\vec{x}$ for all \vec{x} on the unit circle (i.e., with norm 1). We will analyze the behavior of $A\vec{x}$ by using the behavior of $V^\top \vec{x}$, $\Sigma V^\top \vec{x}$ and finally $U\Sigma V^\top \vec{x}$. In the end, we will interpret U as a *rotation or reflection*, Σ as a *scaling*, and V^\top as another *rotation or reflection*.

Before we start, let us discuss what different types of matrices look like as linear transformations. Consider our friendly unit circle:

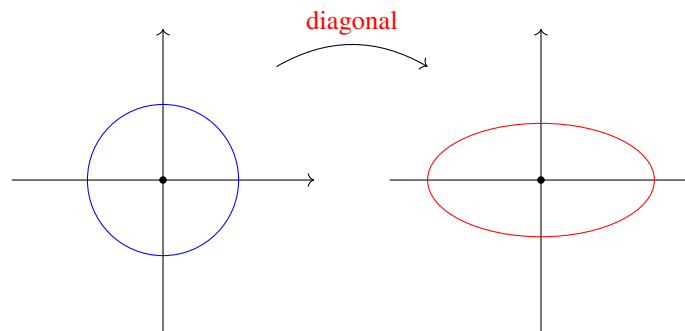


This unit circle consists of all vectors in \mathbb{R}^2 with norm 1.


Multiplying each vector in that circle by the same orthonormal matrix will not change the norm of any vector in that circle, and thus will amount to nothing more than a rotation of the circle, thus not changing its shape.



On the other hand, multiplying each vector in the unit circle by the same diagonal matrix will scale the vectors in the coordinate directions. This means that the unit circle will be mapped, in general, to an axis-aligned ellipse.




In general, matrices aren't orthonormal or diagonal, and so they will both rotate and scale in various ways. This means that the unit circle will be mapped to an ellipse which isn't necessarily axis-aligned.




Let us now systematically study such $A = U\Sigma V^T$ through the lens of the unit circle, as well as where the right singular vectors \vec{v}_1 and \vec{v}_2 are mapped by A .


Since V^T is an orthonormal matrix, it represents a rotation and/or reflection, and so it maps the unit circle to the unit circle, much like our observed figure. Specifically, the right singular vectors \vec{v}_i have $V^T \vec{v}_i = \vec{e}_i$, so they get mapped onto the standard basis by V^T . This gives the following picture.




Now, the diagonal matrix Σ will scale each \vec{e}_i by σ_i , obtaining an ellipse.



Finally, the orthonormal matrix U will map this axis-aligned ellipse to an ellipse which isn't necessarily axis-aligned. Specifically, the vectors that we're looking at, i.e., $\sigma_i \vec{e}_i$, have $U(\sigma_i \vec{e}_i) = \sigma_i U \vec{e}_i = \sigma_i \vec{u}_i$. These \vec{u}_i will be the axes of the resulting ellipse in the same sense as $\sigma_i \vec{e}_i$ were the axes of the axis-aligned ellipse.



Recall that we originally started with a depiction that didn't have any fine-grained description of any vectors, yet obtained the same result:



To understand the impact of A on any general vector \vec{x} , we write it in the V basis: $\vec{x} = \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2$, and use linearity to obtain $A\vec{x} = \alpha_1 \sigma_1 \vec{u}_1 + \alpha_2 \sigma_2 \vec{u}_2$. One can draw this graphically using scaled versions of the above ellipses.

This perspective also says that σ_1 is the maximum scaling of any vector obtained by multiplication by A , and σ_r is the minimum nonzero scaling. (If $r < n$, i.e., A is not full column rank, then there are some nonzero vectors in \mathbb{R}^n which are sent to $\vec{0}$ by A , so the minimum scaling is 0.) You will formally prove this in homework.

2.7 Low-Rank Approximation

Sometimes, in real datasets, matrices have billions or trillions of entries. Storing all of them would be prohibitively expensive, and we would need a way to compress them down into their most important parts. This turns out to be doable via the SVD, as we will see.

To formally talk about a compression algorithm that stores a compressed version of the data with minimal error, we need to talk about what kind of errors are appropriate to discuss in the context of matrices. In the case of vectors, we can use the ℓ^2 -norm, or more generally the ℓ^p norm, to define a distance function; then, the error would just be the distance between the true and the perturbed vectors. This motivates thinking about matrix norms, which allow us to quantify the distance between matrices, and thus create error functions.

There are two ways to think about a matrix. The first way is as a *block of numbers*. Similarly to how we thought of a vector as a block of numbers and found the norm based on this, we can think of the matrix as a big list of vectors and take the norm. This norm is called the *Frobenius norm*, and it corresponds to unrolling an $m \times n$ matrix into a length $m \cdot n$ vector and taking its ℓ^2 -norm.

Definition 37 (Frobenius Norm)

For a matrix $A \in \mathbb{R}^{m \times n}$, its Frobenius norm is defined as

$$\|A\|_F \doteq \sqrt{\sum_{i=1}^m \sum_{j=1}^n A_{ij}^2}. \quad (2.182)$$

The following property will help us when we study vector calculus.

Proposition 38

For a matrix $A \in \mathbb{R}^{m \times n}$, we have $\|A\|_F^2 = \text{tr}(A^\top A)$.

The following pair of results will help us now.

Proposition 39

For a matrix $A \in \mathbb{R}^{m \times n}$ and orthonormal matrices $U \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{n \times n}$, we have

$$\|UAV\|_F = \|UA\|_F = \|AV\|_F = \|A\|_F. \quad (2.183)$$

We do not prove this property now, though it might be on homework.

Proposition 40

For a matrix $A \in \mathbb{R}^{m \times n}$ with rank r and singular values $\sigma_1 \geq \dots \geq \sigma_r > 0$, we have

$$\|A\|_F^2 = \sum_{i=1}^r \sigma_i^2. \quad (2.184)$$

Proof. Let $A = U\Sigma V^\top$ be the SVD of A ; then, we use the previous proposition to get

$$\|A\|_F^2 = \|U\Sigma V^\top\|_F^2 \quad (2.185)$$

$$= \|\Sigma\|_F^2 \quad (2.186)$$

$$= \sum_{i=1}^r \sigma_i^2. \quad (2.187)$$

□

Under this perspective, $\|A - B\|_F$ is small if each component of A and B is close; that is, they are very similar in terms of the block-of-numbers interpretation.

The second way to think about a matrix is as a *linear transformation*. In this case, the matrix is defined by how it acts on vectors via multiplication. A suitable notion of size in this case is the largest scaling factor of the matrix on any unit vector; this is called the *spectral norm* or the *matrix ℓ^2 -norm*.

Definition 41 (Spectral Norm)

For a matrix $A \in \mathbb{R}^{m \times n}$, its spectral norm is defined by

$$\|A\|_2 \doteq \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \|A\vec{x}\|_2. \quad (2.188)$$

Fortunately, this optimization problem has a solution — what’s more, we’ve actually seen this solution before.

Proposition 42

For a matrix $A \in \mathbb{R}^{m \times n}$ with rank r and singular values $\sigma_1 \geq \dots \geq \sigma_r > 0$, we have

$$\|A\|_2 = \sigma_1. \quad (2.189)$$

Proof. We use the Rayleigh quotient to say that

$$\|A\|_2 \doteq \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \|A\vec{x}\|_2 \quad (2.190)$$

$$= \max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \sqrt{\|A\vec{x}\|_2^2} \quad (2.191)$$

$$= \sqrt{\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \|A\vec{x}\|_2^2} \quad (2.192)$$

$$= \sqrt{\max_{\substack{\vec{x} \in \mathbb{R}^n \\ \|\vec{x}\|_2=1}} \vec{x}^\top A^\top A \vec{x}} \quad (2.193)$$

$$= \sqrt{\lambda_{\max}\{A^\top A\}} \quad (2.194)$$

$$= \sigma_1. \quad (2.195)$$

□

Connecting back to the ellipse transformations, the spectral norm captures how much the ellipse is stretched in its most stretched direction. Thus, the matrix is viewed as as a linear map.

To present our main theorems about how to approximate the matrix well under these norms, we need to define notation. Fix a matrix $A \in \mathbb{R}^{m \times n}$. For convenience, let $p \doteq \min\{m, n\}$. Suppose that A has rank $r \leq p$, and that A has SVD

$$A = \sum_{i=1}^p \sigma_i \vec{u}_i \vec{v}_i^\top \quad (2.196)$$

where we note that $\sigma_1 \geq \dots \geq \sigma_r$ and define $\sigma_{r+1} = \sigma_{r+2} = \dots = 0$. Then, for $k \leq p$, we can define

$$A_k \doteq \sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^\top. \quad (2.197)$$

Note that if $k \ll p$, then A_k can be stored much more efficiently than A . For instance, A_k needs to store k scalars (σ s), k vectors in \mathbb{R}^m (\vec{u} s), and k vectors in \mathbb{R}^n (\vec{v} s), for a total storage of $k(m + n + 1)$ floats. On the other hand, storing A requires mn floats naively, and even storing the (full) SVD of A is not much better. So the former is much more efficient to store. The only thing left to do is to show that it actually is a good approximation in A (otherwise what would be the point to storing it instead of A ?)

It turns out that A_k indeed well-approximates A in the sense of the two norms — the Frobenius norm and the spectral norm — that we discussed previously. The two results are collectively known as the *Eckart-Young* (sometimes *Eckart-Young-Mirsky*) theorem(s). We state and prove these now.

We begin with the Eckart-Young theorem for the spectral norm, since it will help us prove the analogous result for Frobenius norms.

Theorem 43 (Eckart-Young Theorem for Spectral Norm)

We have

$$A_k \in \underset{\substack{B \in \mathbb{R}^{m \times n} \\ \text{rank}(B) \leq k}}{\text{argmin}} \|A - B\|_2, \quad (2.198)$$

or, equivalently,

$$\|A - A_k\|_2 \leq \|A - B\|_2, \quad \forall B \in \mathbb{R}^{m \times n}: \text{rank}(B) \leq k. \quad (2.199)$$

Proof of Theorem 43. The proof of [Eckart-Young Theorem for Spectral Norm](#) is partitioned into two parts which together jointly show the conclusion. First, we explicitly calculate $\|A - A_k\|_2$; then for an arbitrary B of rank $\leq k$ we show that $\|A - B\|_2$ is lower-bounded by this quantity.

Part 1. First, we calculate $\|A - A_k\|_2$. Indeed, we have

$$A - A_k = \left(\sum_{i=1}^p \sigma_i \vec{u}_i \vec{v}_i^\top \right) - \left(\sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^\top \right) \quad (2.200)$$

$$= \sum_{i=k+1}^p \sigma_i \vec{u}_i \vec{v}_i^\top. \quad (2.201)$$

Since $\sigma_{k+1} \geq \sigma_{k+2} \geq \dots$, the \vec{u}_i are orthonormal, and the \vec{v}_i are orthonormal, this is a valid singular value decomposition of the rank $r - k$ matrix $A - A_k$. Thus $A - A_k$ has largest singular value equal to σ_{k+1} , so

$$\|A - A_k\|_2 = \sigma_{k+1}. \quad (2.202)$$

Part 2. We now show that for any matrix $B \in \mathbb{R}^{m \times n}$ of rank $\leq k$, we have $\|A - B\|_2 \geq \sigma_{k+1}$. Before we commence with the proof, let us first discuss the overall argument structure.

We have two characterizations of the spectral norm: the maximum singular value, and the maximal Rayleigh coefficient. Up until now, we don't have any advanced machinery such as singular value inequalities to directly show that the maximum singular value of $A - B$ is lower-bounded by some constant. However, we do understand matrix-vector products, and how to characterize their optima, pretty well, so this motivates the use of the maximal Rayleigh coefficient. In particular, suppose $f(\vec{x}) = \|(A - B)\vec{x}\|_2 / \|\vec{x}\|_2$. Then we know that for any specific value of \vec{x}_0 , we have $\|A - B\|_2 = \max_{\vec{x}} f(\vec{x}) \geq f(\vec{x}_0)$. Thus, one way to get a lower bound on $\|A - B\|_2$ is to plug in any \vec{x}_0 into f . The trick is, given B , to find the right value of \vec{x}_0 such that $f(\vec{x}_0) = \sigma_{k+1}$. The first part of the argument will be finding an appropriate \vec{x}_0 ; the next will show that it works.

Okay, let's start with the proof.

Step 1. Here we show the existence of an appropriate choice for \vec{x}_0 .

Because B has rank $\leq k$, the rank-nullity theorem gives

$$\dim(\mathcal{N}(B)) = n - \text{rank}(B) \geq n - k > 0. \quad (2.203)$$

Now, define $V_{k+1} \doteq [\vec{v}_1, \dots, \vec{v}_{k+1}]$. We know by the SVD that the \vec{v}_i are orthonormal, and so are linearly independent. Thus

$$\text{rank}(V_{k+1}) = \dim(\mathcal{R}(V_{k+1})) = k + 1. \quad (2.204)$$

Thus

$$\dim(\mathcal{N}(B)) + \dim(\mathcal{R}(V_{k+1})) \geq (n - k) + (k + 1) = n + 1 > n. \quad (2.205)$$

Claim. There exists a unit vector in $\mathcal{N}(B) \cap \mathcal{R}(V_{k+1})$.

Proof. First, we note that since $\mathcal{N}(B)$ and $\mathcal{R}(V_{k+1})$ are subspaces of \mathbb{R}^n , their intersection is a subspace of \mathbb{R}^n . Thus, it is closed under scalar multiplication. Thus the existence of a unit vector is equivalent to the existence of a nonzero vector, since one can scale this nonzero vector by its (inverse) norm to get the unit vector.

Suppose for the sake of contradiction that $\mathcal{N}(B)$ and $\mathcal{R}(V_{k+1})$ have no nonzero vectors in common. Fix a basis S_1 for $\mathcal{N}(B)$ and a basis S_2 for $\mathcal{R}(V_{k+1})$. By assumption, we have that $S_1 \cap S_2 = \emptyset$, and $S_1 \cup S_2$ is a linearly independent set. Here, there are two cases.

- A. $S_1 \cup S_2$ contains $< n + 1$ vectors. This means that S_1 and S_2 have a vector in common, so $S_1 \cap S_2 \neq \emptyset$ and we have reached a contradiction.
- B. $S_1 \cup S_2$ has exactly $n + 1$ vectors. As a set of $n + 1$ vectors in \mathbb{R}^n , $S_1 \cup S_2$ is linearly dependent and we have reached a contradiction.

In both cases we have reached a contradiction, so there is a nonzero vector in $\mathcal{N}(B) \cap \mathcal{R}(V_{k+1})$.

We may scale this vector by its inverse norm to get a unit vector in $\mathcal{N}(B) \cap \mathcal{R}(V_{k+1})$. \square

Let $\vec{x}_0 \in \mathcal{N}(B) \cap \mathcal{R}(V_{k+1})$ be any such unit vector. We use this vector for the Rayleigh coefficient.

Step 2. Now we show that our choice of \vec{x}_0 plugged into the Rayleigh coefficient gives σ_{k+1} as a lower bound.

Indeed, we have

$$\|A - B\|_2 = \max_{\vec{x} \neq \vec{0}} \frac{\|(A - B)\vec{x}\|_2}{\|\vec{x}\|_2} \quad (2.206)$$

$$\geq \frac{\|(A - B)\vec{x}_0\|_2}{\|\vec{x}_0\|_2} \quad (2.207)$$

$$= \|(A - B)\vec{x}_0\|_2 \quad (2.208)$$

$$= \|A\vec{x}_0 - B\vec{x}_0\|_2. \quad (2.209)$$

Since $\vec{x}_0 \in \mathcal{N}(B)$, we have $B\vec{x}_0 = \vec{0}$, so

$$\|A - B\|_2 \geq \|A\vec{x}_0 - B\vec{x}_0\|_2 = \|A\vec{x}_0\|_2. \quad (2.210)$$

Since $\vec{x}_0 \in \mathcal{R}(V_{k+1})$, there are constants $\alpha_1, \dots, \alpha_{k+1}$, not all zero, such that $\vec{x}_0 = \sum_{i=1}^{k+1} \alpha_i \vec{v}_i$. Thus

$$\|A - B\|_2 \geq \|A\vec{x}_0\|_2 \quad (2.211)$$

$$= \left\| A \sum_{i=1}^{k+1} \alpha_i \vec{v}_i \right\|_2 \quad (2.212)$$

$$= \left\| \left(\sum_{i=1}^p \sigma_i \vec{u}_i \vec{v}_i^\top \right) \left(\sum_{i=1}^{k+1} \alpha_i \vec{v}_i \right) \right\|_2 \quad (2.213)$$

$$= \left\| \sum_{i=1}^p \sum_{j=1}^{k+1} \sigma_i \alpha_j \vec{u}_i \vec{v}_i^\top \vec{v}_j \right\|_2. \quad (2.214)$$

By vector algebra, the fact that the \vec{u}_i are orthonormal, and the fact that the \vec{v}_i are orthonormal, one can mechanically show that

$$\|A - B\|_2 \geq \left\| \sum_{i=1}^p \sum_{j=1}^{k+1} \sigma_i \alpha_j \vec{u}_i \vec{v}_i^\top \vec{v}_j \right\|_2 \quad (2.215)$$

$$= \left\| \sum_{i=1}^{k+1} \alpha_i \sigma_i \vec{u}_i \right\|_2 \quad (2.216)$$

$$= \sqrt{\left\| \sum_{i=1}^{k+1} \alpha_i \sigma_i \vec{u}_i \right\|_2^2} \quad (2.217)$$

$$= \sqrt{\sum_{i=1}^{k+1} \alpha_i^2 \sigma_i^2} \quad (2.218)$$

$$\geq \sqrt{\sigma_{k+1}^2 \sum_{i=1}^{k+1} \alpha_i^2} \quad (2.219)$$

$$= \sigma_{k+1} \sqrt{\sum_{i=1}^{k+1} \alpha_i^2} \quad (2.220)$$

$$= \sigma_{k+1} \|\vec{x}_0\|_2 \quad (2.221)$$

$$= \sigma_{k+1}. \quad (2.222)$$

This proves the claim. □

We now state and prove the other result, which is the Eckart-Young theorem for the Frobenius norm.

Theorem 44 (Eckart-Young Theorem for Frobenius Norm)

We have

$$A_k \in \underset{\substack{B \in \mathbb{R}^{m \times n} \\ \text{rank}(B) \leq k}}{\text{argmin}} \|A - B\|_F, \quad (2.223)$$

or, equivalently,

$$\|A - A_k\|_F^2 \leq \|A - B\|_F^2, \quad \forall B \in \mathbb{R}^{m \times n} : \text{rank}(B) \leq k. \quad (2.224)$$

Proof of Theorem 44. Like the Frobenius norm, the proof of [Eckart-Young Theorem for Frobenius Norm](#) is partitioned into two parts which together jointly show the conclusion. First, we explicitly calculate $\|A - A_k\|_F^2$; then for an arbitrary B of rank $\leq k$ we show that $\|A - B\|_F^2$ is lower-bounded by this quantity.

Part 1. First, we calculate $\|A - A_k\|_F^2$. Indeed, we have

$$A - A_k = \left(\sum_{i=1}^p \sigma_i \vec{u}_i \vec{v}_i^\top \right) - \left(\sum_{i=1}^k \sigma_i \vec{u}_i \vec{v}_i^\top \right) \quad (2.225)$$

$$= \sum_{i=k+1}^p \sigma_i \vec{u}_i \vec{v}_i^\top. \quad (2.226)$$

Since $\sigma_{k+1} \geq \sigma_{k+2} \geq \dots$, the \vec{u}_i are orthonormal, and the \vec{v}_i are orthonormal, this is a valid singular value decomposition of the rank $r - k$ matrix $A - A_k$. Thus $A - A_k$ has largest singular value equal to σ_{k+1} , so

$$\|A - A_k\|_F^2 = \sum_{i=k+1}^p \sigma_i^2. \quad (2.227)$$

Part 2. We now show that for any matrix $B \in \mathbb{R}^{m \times n}$ of rank $\leq k$, we have $\|A - B\|_F^2 \geq \sum_{i=k+1}^p \sigma_i^2$. Again, let us begin by discussing the overall argument structure.

Our goal is to show an inequality between two sums of squares of singular values. To do this, it is simplest to match up terms in the two sums and show the inequality between each pair of terms. Then summing over all terms preserves the inequality.

More concretely, define $C \doteq A - B$. Suppose C has SVD $C = \sum_{i=1}^p \gamma_i \vec{y}_i \vec{z}_i^\top$. We want to show

$$\|A - B\|_F^2 = \sum_{i=1}^p \gamma_i^2 \underset{\text{still need to show}}{\geq} \sum_{i=k+1}^p \sigma_i^2. \quad (2.228)$$

Thus, we claim that $\gamma_i \geq \sigma_{i+k}$ for every $i \in \{1, \dots, p\}$, with the understanding that $\sigma_{p+1} = \sigma_{p+2} = \dots = 0$.

To use our earlier knowledge about the spectral norm to our advantage, we write the singular value γ_i as the spectral norm of the approximation error matrix:

$$\gamma_i = \|C - C_{i-1}\|_2. \quad (2.229)$$

We can expand to get

$$\gamma_i = \|C - C_{i-1}\|_2 = \|(A - B) - C_{i-1}\|_2 = \|A - (B + C_{i-1})\|_2. \quad (2.230)$$

Now this looks somewhat like the low-rank approximation setup, because we are computing the spectral norm of the difference between A and some matrix. But what is the rank of this matrix? We have

$$\text{rank}(B + C_{i-1}) \leq \text{rank}(B) + \text{rank}(C_{i-1}) \leq k + i - 1. \quad (2.231)$$

Thus, by applying [Eckart-Young Theorem for Spectral Norm](#) to the rank $i + k - 1$ approximation of A , that

$$\gamma_i = \|A - (B + C_{i-1})\|_2 \geq \|A - A_{i+k-1}\|_2 = \sigma_{i+k} \quad (2.232)$$

which is what we wanted to show.

To finish the proof, we use the following inequalities:

$$\gamma_i^2 \geq \sigma_{i+k}^2, \quad \forall i \in \{1, \dots, p\} \quad (2.233)$$

$$\sum_{i=1}^p \gamma_i^2 \geq \sum_{i=1}^p \sigma_{i+k}^2 \quad \text{summing over all } i, \quad (2.234)$$

$$\|A - B\|_F^2 \geq \sum_{i=k+1}^p \sigma_i^2 \quad (2.235)$$

as desired. □

2.8 (OPTIONAL) Block Matrix Identities

In this section, we list many ways to manipulate block matrices. Since each fact in here is something you can derive yourself using definitions (e.g. of matrix multiplication), you may use any of them without proof.

2.8.1 Transposes of Block Matrices

$$\begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix}^\top = \begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} \quad (2.236)$$

$$\begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} = \begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix}^\top \quad (2.237)$$

$$\begin{bmatrix} A & B \end{bmatrix}^\top = \begin{bmatrix} A^\top \\ B^\top \end{bmatrix} \quad (2.238)$$

$$\begin{bmatrix} A \\ B \end{bmatrix}^\top = \begin{bmatrix} A^\top & B^\top \end{bmatrix} \quad (2.239)$$

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^\top = \begin{bmatrix} A^\top & C^\top \\ B^\top & D^\top \end{bmatrix}. \quad (2.240)$$

2.8.2 Block Matrix Products

In the following, \vec{e}_i is the i^{th} standard basis vector – it has a 1 in the i^{th} coordinate and 0 in all other coordinates.

$$\begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix} \begin{bmatrix} \vec{y}_1^\top \\ \vdots \\ \vec{y}_n^\top \end{bmatrix} = \sum_{i=1}^n \vec{x}_i \vec{y}_i^\top \quad (2.241)$$

$$\begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} \begin{bmatrix} \vec{y}_1 & \cdots & \vec{y}_n \end{bmatrix} = \begin{bmatrix} \vec{x}_1^\top \vec{y}_1 & \cdots & \vec{x}_1^\top \vec{y}_n \\ \vdots & \ddots & \vdots \\ \vec{x}_n^\top \vec{y}_1 & \cdots & \vec{x}_n^\top \vec{y}_n \end{bmatrix} \quad (2.242)$$

$$\vec{e}_i^\top \begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} = \vec{x}_i^\top \quad (2.243)$$

$$\begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix} \vec{e}_i = \vec{x}_i \quad (2.244)$$

$$A \begin{bmatrix} \vec{x}_1 & \cdots & \vec{x}_n \end{bmatrix} = \begin{bmatrix} A\vec{x}_1 & \cdots & A\vec{x}_n \end{bmatrix} \quad (2.245)$$

$$\begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} A = \begin{bmatrix} \vec{x}_1^\top A \\ \vdots \\ \vec{x}_n^\top A \end{bmatrix} \quad (2.246)$$

$$A \begin{bmatrix} B & C \end{bmatrix} = \begin{bmatrix} AB & AC \end{bmatrix} \quad (2.247)$$

$$\begin{bmatrix} A \\ B \end{bmatrix} C = \begin{bmatrix} AC \\ BC \end{bmatrix}. \quad (2.248)$$

2.8.3 Block Diagonal Matrices

$$\begin{bmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_n \end{bmatrix} \begin{bmatrix} \vec{x}_1^\top \\ \vec{x}_2^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} = \begin{bmatrix} d_1 \vec{x}_1^\top \\ \vdots \\ d_n \vec{x}_n^\top \end{bmatrix} \quad (2.249)$$

$$\begin{bmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ \vdots \\ B_n \end{bmatrix} = \begin{bmatrix} A_1 B_1 \\ A_2 B_2 \\ \vdots \\ A_n B_n \end{bmatrix} \quad (2.250)$$

2.8.4 Quadratic Forms

$$\vec{x}^\top A \vec{y} = \sum_i \sum_j A_{ij} x_i y_j \quad (2.251)$$

$$\begin{bmatrix} \vec{x} \\ \vec{y} \end{bmatrix}^\top \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} \vec{x} \\ \vec{y} \end{bmatrix} = \vec{x}^\top A \vec{x} + \vec{x}^\top B \vec{y} + \vec{y}^\top C \vec{x} + \vec{y}^\top D \vec{y}. \quad (2.252)$$

Chapter 3

Vector Calculus

Relevant sections of the textbooks:

- [1] Appendix A.4.

3.1 Gradient, Jacobian, and Hessian

To motivate this section, we start with a familiar concept which should have been covered in the prerequisites: the derivatives of a *scalar* function $f: \mathbb{R} \rightarrow \mathbb{R}$ which takes in scalar input and produces a scalar output. The derivative quantifies the (instantaneous) rate of change of the function due to the change of its input. We recall the limit definition of the derivative.

Definition 45 (Derivative for Scalar Functions)

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be differentiable. The *derivative* of f with respect to x is the function $\frac{df}{dx}: \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$\frac{df}{dx}(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}. \quad (3.1)$$

The derivative of f has several other notations, including f' and \dot{f} .

In this section, we aim to generalize the concept of derivatives beyond scalar functions. We will focus on two types of functions:

1. *Multivariate functions* $f: \mathbb{R}^n \rightarrow \mathbb{R}$ which take a vector $\vec{x} \in \mathbb{R}^n$ as input and produce a scalar $f(\vec{x}) \in \mathbb{R}$ as output. Familiar examples of such functions include $f(\vec{x}) = \|\vec{x}\|_p$ and $f(\vec{x}) = \vec{a}^\top \vec{x}$.
2. *Vector-valued functions* $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ which take a vector $\vec{x} \in \mathbb{R}^n$ as input and produce another vector $\vec{f}(\vec{x}) \in \mathbb{R}^m$ as output. A familiar example of such functions is $f(\vec{x}) = A\vec{x}$.

One tool that allows us to compute derivatives of scalar functions is the chain rule, which describes the derivative of the composition of two functions.

Theorem 46 (Chain Rule for Scalar Functions)

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ be two differentiable scalar functions, and let $h: \mathbb{R} \rightarrow \mathbb{R}$ be defined as $h(x) =$

$f(g(x))$ for all $x \in \mathbb{R}$. Then h is differentiable, and

$$\frac{dh}{dx}(x) = \frac{df}{dg}(g(x)) \cdot \frac{dg}{dx}(x). \quad (3.2)$$

Here, we use some — perhaps unfamiliar — notation not previously introduced. In particular, the derivative $\frac{df}{dg}(g(x))$ is a little strange. When we write $\frac{df}{dg}$ we take the derivative of f with respect to the output of g . In this case, we know that the input of f is exactly the output of g , so really $\frac{df}{dg}$ takes the derivative of f with respect to its input. Thus we can more compactly write the chain rule using the f' notation as

$$h'(x) = f'(g(x)) \cdot g'(x). \quad (3.3)$$

We will see in the next section that the chain rule can be generalized to settings of multivariate functions. To aid our study of such generalizations, we will introduce here a *computational graph* perspective of the chain rule. At the moment this graphical perspective looks trivial, but it will help us understand more complicated cases.




Figure 3.1: Graphical depiction of the single-variable chain rule corresponding to $h = f \circ g$.

In this computational graph, one computes the derivative by summing along all paths from x to $h(x)$. There is only one path, so we get

$$\frac{dh}{dx}(x) = \frac{df}{dg}(g(x)) \cdot \frac{dg}{dx}(x). \quad (3.4)$$

3.1.1 Partial Derivatives

For multivariate functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$, when we talk about the rate of change of the function with respect to its input, we need to specify which input we are talking about. Partial derivatives quantify this and give us the rate of change of the function due to the change of one of its inputs, say x_i , while keeping all other inputs fixed. Keeping all but one input fixed renders a scalar function, for which we know how to compute the derivative. To formalize this we introduce the limit definition of the partial derivative.

Definition 47 (Partial Derivative)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable. The *partial derivative* of f with respect to x_i is the function $\frac{\partial f}{\partial x_i}: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$\frac{\partial f}{\partial x_i}(\vec{x}) = \lim_{h \rightarrow 0} \frac{f(x_1, \dots, x_i + h, \dots, x_n) - f(\vec{x})}{h}, \quad (3.5)$$

or equivalently,

$$\frac{\partial f}{\partial x_i}(\vec{x}) = \lim_{h \rightarrow 0} \frac{f(\vec{x} + h \cdot \vec{e}_i) - f(\vec{x})}{h} \quad (3.6)$$

where \vec{e}_i is the i^{th} standard basis vector.

This limit definition gives an alternative way of interpreting partial derivatives: $\frac{\partial f}{\partial x_i}$ is the rate of change of the function along the direction of the standard basis vector \vec{e}_i .

In practice, we do not use the limit definition to compute regular derivatives. Similarly, we do not use the limit definition to compute partial derivatives. The main way to compute partial derivatives uses the following tip.

Problem Solving Strategy 48. To compute the partial derivative $\frac{\partial f}{\partial x_i}$, pretend that all x_j for $j \neq i$ are constants, then take the ordinary derivative in x_i .

Example 49. Consider the function $f(\vec{x}) = \vec{a}^\top \vec{x}$. Then

$$\frac{\partial f}{\partial x_i}(\vec{x}) = \frac{\partial}{\partial x_i} \vec{a}^\top \vec{x} \quad (3.7)$$

$$= \frac{\partial}{\partial x_i} \sum_{j=1}^n a_j x_j \quad (3.8)$$

$$= \frac{\partial}{\partial x_i} (a_1 x_1 + \cdots + a_n x_n) \quad (3.9)$$

$$= a_i. \quad (3.10)$$

Let us consider the case where the input \vec{x} is not an independent vector but rather depends on another variable $t \in \mathbb{R}$, i.e., $\vec{x}: \mathbb{R} \rightarrow \mathbb{R}^n$ is a function. In such case the function $f(\vec{x}(t))$ has *one independent* input, which is t . If we are interested in finding the derivative of f with respect to t , we can utilize a chain rule to do so.

Theorem 50 (Chain Rule For Multivariate Functions)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and $\vec{g}: \mathbb{R} \rightarrow \mathbb{R}^n$ be differentiable functions. Define the function $h: \mathbb{R} \rightarrow \mathbb{R}$ by $h(x) = f(\vec{g}(x))$ for all $x \in \mathbb{R}$. Then h is differentiable and has derivative

$$\frac{dh}{dx}(x) = \sum_{i=1}^n \frac{\partial f}{\partial g_i}(\vec{g}(x)) \cdot \frac{dg_i}{dx}(x). \quad (3.11)$$

Here we again review this computational graph perspective. There are now n separate paths from x to h , like so:




Figure 3.2: Graphical depiction of the multivariate chain rule corresponding to $h = f \circ \vec{g}$.

In this computational graph, one computes the derivative by summing across the n paths from x to $h(x)$, obtaining

$$\frac{dh}{dx}(x) = \sum_{i=1}^n \frac{\partial f}{\partial g_i}(\vec{g}(x)) \cdot \frac{dg_i}{dx}(x). \quad (3.12)$$

Here again we use the notation $\frac{\partial f}{\partial g_i}$ to denote the derivative of f with respect to the output of g_i . Since the output of g_i is really the i^{th} input of f , this derivative is just the derivative of f with respect to its i^{th} input.

3.1.2 Gradient

We will now use the definition of partial derivatives to introduce the gradient of multivariate functions.

Definition 51 (Gradient)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function. The *gradient* of f is the function $\nabla f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by

$$\nabla f(\vec{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\vec{x}) \\ \vdots \\ \frac{\partial f}{\partial x_n}(\vec{x}) \end{bmatrix}. \quad (3.13)$$

Note that the gradient is a column vector. The transpose of the gradient is (confusingly!) referred to as the *derivative* of the function. We will now list two important geometric properties of the gradient. The first can be stated straight away:

Proposition 52

Let $\vec{x} \in \mathbb{R}^n$. The gradient $\nabla f(\vec{x})$ points in the direction of *steepest ascent* at \vec{x} , i.e., the direction around \vec{x} in which f has the maximum rate of change. Furthermore, this rate of change is quantified by the norm $\|\nabla f(\vec{x})\|_2$.

Proof. Let $\vec{u} \in \mathbb{R}^n$ be a unit vector (i.e., representing an arbitrary direction in \mathbb{R}^n). Using the Cauchy-Schwarz inequality we can write:

$$\vec{u}^\top [\nabla f(\vec{x})] \leq \|\vec{u}\|_2 \|\nabla f(\vec{x})\|_2 \quad (3.14)$$

$$= \|\nabla f(\vec{x})\|_2, \quad (3.15)$$

so the maximum value that the expression can take is $\|\nabla f(\vec{x})\|_2$. Now it remains to show that this value is attained for the choice $\vec{u} = \frac{\nabla f(\vec{x})}{\|\nabla f(\vec{x})\|_2}$.

$$\left[\frac{\nabla f(\vec{x})}{\|\nabla f(\vec{x})\|_2} \right]^\top [\nabla f(\vec{x})] = \frac{\|\nabla f(\vec{x})\|_2^2}{\|\nabla f(\vec{x})\|_2} \quad (3.16)$$

$$= \|\nabla f(\vec{x})\|_2. \quad (3.17)$$

□

To list the second property, first we need a quick definition.

Definition 53 (Level Set)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function, and $\alpha \in \mathbb{R}$ be a scalar.

- The α -*level set* of f is the set of points \vec{x} such that $f(\vec{x}) = \alpha$:

$$L_\alpha(f) = \{\vec{x} \in \mathbb{R}^n \mid f(\vec{x}) = \alpha\}. \quad (3.18)$$

- The α -*sublevel set* of f is the set of points \vec{x} such that $f(\vec{x}) \leq \alpha$:

$$L_{\leq \alpha}(f) = \{\vec{x} \in \mathbb{R}^n \mid f(\vec{x}) \leq \alpha\}. \quad (3.19)$$

- The α -superlevel set of f is the set of points \vec{x} such that $f(\vec{x}) \geq \alpha$:

$$L_{\geq\alpha}(f) = \{\vec{x} \in \mathbb{R}^n \mid f(\vec{x}) \geq \alpha\}. \quad (3.20)$$

Proposition 54

Let $\vec{x} \in \mathbb{R}^n$ and suppose $f(\vec{x}) = \alpha$. Then $\nabla f(\vec{x})$ is orthogonal to the hyperplane which is tangent at \vec{x} to the α -level set of f .

We illustrate the two properties through examples.

Example 55 (Gradient of the Squared ℓ^2 Norm). In this example we will compute and visualize the gradient of the function $f(\vec{x}) = \|\vec{x}\|_2^2$ where $\vec{x} \in \mathbb{R}^2$.

$$\nabla f(\vec{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\vec{x}) \\ \frac{\partial f}{\partial x_2}(\vec{x}) \end{bmatrix} \quad (3.21)$$

$$= \begin{bmatrix} \frac{\partial}{\partial x_1}(x_1^2 + x_2^2) \\ \frac{\partial}{\partial x_2}(x_1^2 + x_2^2) \end{bmatrix} \quad (3.22)$$

$$= \begin{bmatrix} 2x_1 \\ 2x_2 \end{bmatrix} \quad (3.23)$$

$$= 2\vec{x}. \quad (3.24)$$

This function has a paraboloid-shaped graph, as shown in Figure 3.3a. Let us now find the α -level set of the function f for some constant α .

$$\alpha = f(\vec{x}) \quad (3.25)$$

$$= x_1^2 + x_2^2 \quad (3.26)$$

For a given $\alpha \geq 0$, the α -level set is a circle centered at the origin which has radius $\sqrt{\alpha}$. Now we evaluate the gradient at a few points on these level sets:

$$\nabla f(-1, 0) = (-2, 0) \quad (3.27)$$

$$\nabla f(1/\sqrt{2}, 1/\sqrt{2}) = (\sqrt{2}, \sqrt{2}) \quad (3.28)$$

$$\nabla f(2, 0) = (4, 0) \quad (3.29)$$

$$\nabla f(-\sqrt{2}, \sqrt{2}) = (-2\sqrt{2}, 2\sqrt{2}) \quad (3.30)$$

We plot the level sets for $\alpha = 1$ and $\alpha = 4$ and visualize the gradient directions in Figure 3.3b.

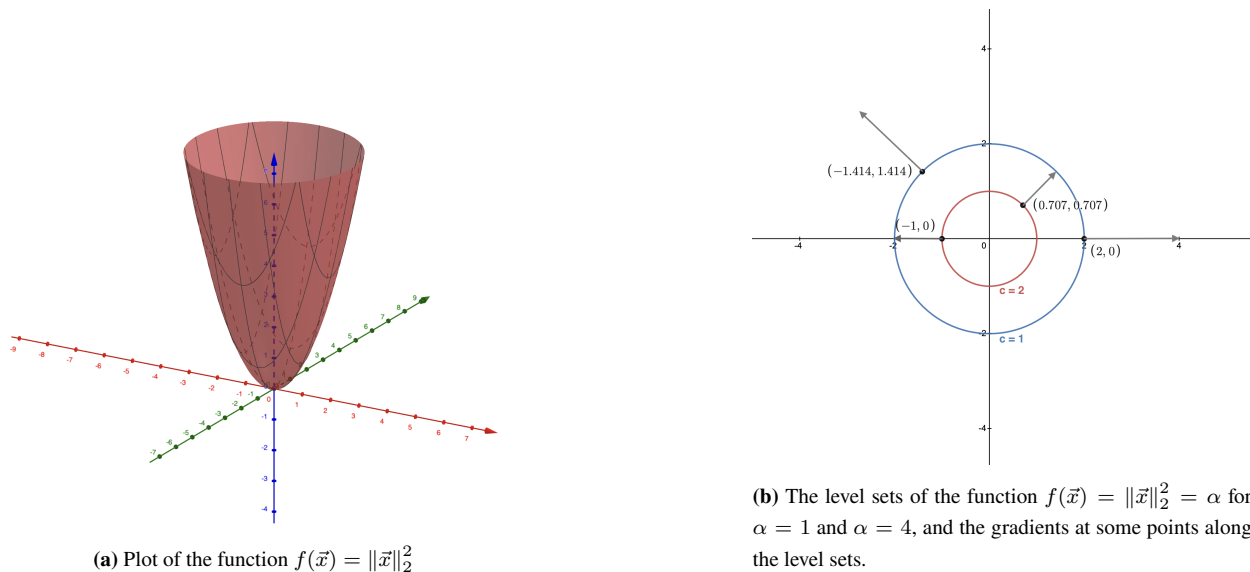


Figure 3.3: Example 55

We make the following observations:

1. At each point on a given level set, the gradient vector at that point is orthogonal to the line tangent to the level set at that point.
2. The length of the gradient vector increases as we move away from the origin. This means that the function gets steeper in that direction (i.e., it changes more rapidly).

Example 56 (Gradient of Linear Function). In this example we will compute and comment on the gradient of the linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by $f(\vec{x}) = \vec{a}^\top \vec{x}$ where $\vec{a} \in \mathbb{R}^n$ is fixed. We have

$$f(\vec{x}) = \vec{a}^\top \vec{x}, \quad (3.31)$$

$$\nabla f(\vec{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\vec{x}) \\ \vdots \\ \frac{\partial f}{\partial x_n}(\vec{x}) \end{bmatrix} \quad (3.32)$$

$$= \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \quad (3.33)$$

$$= \vec{a}. \quad (3.34)$$

We make the following observations:

1. The α -level sets of this function are the hyperplanes given by all \vec{x} such that $\vec{a}^\top \vec{x} = \alpha$. These hyperplanes have normal vector \vec{a} . Notice that the normal vector (which is orthogonal to all vectors in the hyperplane) is exactly the gradient vector.
2. The gradient is constant, meaning that the function has a constant rate of change everywhere.

Example 57 (Gradient of the Quadratic Form). Let $A \in \mathbb{R}^{n \times n}$. In this example we will compute the gradient of the quadratic function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by $f(\vec{x}) = \vec{x}^\top A \vec{x}$. Indeed, we have

$$f(\vec{x}) = \vec{x}^\top A \vec{x} \quad (3.35)$$

$$= \sum_{i=1}^n \sum_{j=1}^n A_{ij} x_i x_j. \quad (3.36)$$

Now fix $k \in \{1, \dots, n\}$, and we find the partial derivative with respect to x_k . We have

$$\begin{aligned} f(\vec{x}) &= \sum_{i=1}^n \sum_{j=1}^n A_{ij} x_i x_j \\ &= \sum_{j=1}^n A_{kj} x_k x_j + \sum_{\substack{i=1 \\ i \neq k}}^n \sum_{j=1}^n A_{ij} x_i x_j \\ &= \sum_{j=1}^n A_{kj} x_k x_j + \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i x_k + \sum_{\substack{i=1 \\ i \neq k}}^n \sum_{\substack{j=1 \\ j \neq k}}^n A_{ij} x_i x_j \\ &= A_{kk} x_k^2 + \sum_{\substack{j=1 \\ j \neq k}}^n A_{kj} x_k x_j + \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i x_k + \sum_{\substack{i=1 \\ i \neq k}}^n \sum_{\substack{j=1 \\ j \neq k}}^n A_{ij} x_i x_j \\ &= A_{kk} x_k^2 + x_k \sum_{\substack{j=1 \\ j \neq k}}^n A_{kj} x_j + x_k \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i + \sum_{\substack{i=1 \\ i \neq k}}^n \sum_{\substack{j=1 \\ j \neq k}}^n A_{ij} x_i x_j. \end{aligned}$$

Then, taking the derivatives, we have

$$\frac{\partial f}{\partial x_k}(\vec{x}) = \frac{\partial}{\partial x_k} \left(A_{kk} x_k^2 + x_k \sum_{\substack{j=1 \\ j \neq k}}^n A_{kj} x_j + x_k \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i + \sum_{\substack{i=1 \\ i \neq k}}^n \sum_{\substack{j=1 \\ j \neq k}}^n A_{ij} x_i x_j \right) \quad (3.37)$$

$$= 2A_{kk} x_k + \sum_{\substack{j=1 \\ j \neq k}}^n A_{kj} x_j + \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i \quad (3.38)$$

$$= \left(A_{kk} x_k + \sum_{\substack{j=1 \\ j \neq k}}^n A_{kj} x_j \right) + \left(A_{kk} x_k + \sum_{\substack{i=1 \\ i \neq k}}^n A_{ik} x_i \right) \quad (3.39)$$

$$= \sum_{j=1}^n A_{kj} x_j + \sum_{i=1}^n A_{ik} x_i \quad (3.40)$$

$$= (A \vec{x})_k + (A^\top \vec{x})_k \quad (3.41)$$

$$= ((A + A^\top) \vec{x})_k. \quad (3.42)$$

Thus computing the gradient via stacking the partial derivatives gets

$$\nabla f(\vec{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\vec{x}) \\ \vdots \\ \frac{\partial f}{\partial x_n}(\vec{x}) \end{bmatrix} \quad (3.43)$$

$$= \begin{bmatrix} ((A + A^\top)\vec{x})_1 \\ \vdots \\ ((A + A^\top)\vec{x})_n \end{bmatrix} \quad (3.44)$$

$$= (A + A^\top)\vec{x}. \quad (3.45)$$

That was a very involved computation! Luckily, in the near future, we will see ways to simplify the process of computing gradients.

3.1.3 Jacobian

We now have the tools to generalize the notion of derivatives to vector-valued functions $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$.

Definition 58 (Jacobian)

Let $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a differentiable function. The *Jacobian* of \vec{f} is the function $D\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ defined as

$$D\vec{f}(\vec{x}) = \begin{bmatrix} \nabla f_1(\vec{x})^\top \\ \vdots \\ \nabla f_m(\vec{x})^\top \end{bmatrix} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1}(\vec{x}) & \cdots & \frac{\partial f_1}{\partial x_n}(\vec{x}) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1}(\vec{x}) & \cdots & \frac{\partial f_m}{\partial x_n}(\vec{x}) \end{bmatrix}. \quad (3.46)$$

One big thing to note is that *the Jacobian is different from the gradient!* If $f: \mathbb{R}^n \rightarrow \mathbb{R}^1 = \mathbb{R}$, then its Jacobian $Df: \mathbb{R}^n \rightarrow \mathbb{R}^{1 \times n}$ is a function which outputs a *row vector*. This row vector is the *transpose* of the gradient.

We can develop a nice and general chain rule with the Jacobian.

Theorem 59 (Chain Rule for Vector-Valued Functions)

Let $\vec{f}: \mathbb{R}^p \rightarrow \mathbb{R}^m$ and $\vec{g}: \mathbb{R}^n \rightarrow \mathbb{R}^p$ be differentiable functions. Let $\vec{h}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be defined as $\vec{h}(\vec{x}) = \vec{f}(\vec{g}(\vec{x}))$ for all $\vec{x} \in \mathbb{R}^n$. Then \vec{h} is differentiable, and

$$D\vec{h}(\vec{x}) = [D\vec{f}(\vec{g}(\vec{x}))][D\vec{g}(\vec{x})]. \quad (3.47)$$

Here, as before, the notation $D\vec{f}(\vec{g}(\vec{x}))$ means that we compute $D\vec{f}$ and then evaluate it on the point $\vec{g}(\vec{x})$.

This is a broad chain rule which we can apply to many problems, but we must always remember that *for a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, its Jacobian is the transpose of its gradient*. One typical chain rule we can derive from the general one follows below.

Corollary 60. Let $f: \mathbb{R}^p \rightarrow \mathbb{R}$ and $\vec{g}: \mathbb{R}^n \rightarrow \mathbb{R}^p$ be differentiable functions. Let $h: \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as $h(\vec{x}) = f(\vec{g}(\vec{x}))$ for all $\vec{x} \in \mathbb{R}^n$. Then h is differentiable, and

$$\nabla h(\vec{x}) = [D\vec{g}(\vec{x})]^\top \nabla f(\vec{g}(\vec{x})). \quad (3.48)$$

Example 61. In this example we will use the chain rule to compute the gradient of the function $h(\vec{x}) = \|A\vec{x} - \vec{y}\|_2^2$. It can be written as $h(\vec{x}) = f(\vec{g}(\vec{x}))$ where $f(\vec{x}) = \|\vec{x}\|_2^2$ and $\vec{g}(\vec{x}) = A\vec{x} - \vec{y}$. We have that

$$D\vec{g}(\vec{x}) = A \quad (3.49)$$

(the proof is in discussion or homework), and also from earlier

$$\nabla f(\vec{x}) = 2\vec{x}. \quad (3.50)$$

Thus applying the chain rule obtains

$$\nabla h(\vec{x}) = [D\vec{g}(\vec{x})]^\top \nabla f(\vec{g}(\vec{x})) \quad (3.51)$$

$$= 2A^\top (A\vec{x} - \vec{y}) \quad (3.52)$$

as desired. This gradient matches what would be computed if we had used the componentwise partial derivatives.

Finally, we again revisit the computational graph perspective. Here, the graph looks like:




Figure 3.4: Graphical depiction of the multivariate chain rule corresponding to $\vec{h} = \vec{f} \circ \vec{g}$.

To find the partial derivative $\frac{\partial f_k}{\partial x_i}(\vec{x})$, one sums across all paths in the graph from x_i to f_k , i.e.,

$$[D\vec{h}(\vec{x})]_{k,i} = \frac{\partial f_k}{\partial x_i}(\vec{x}) \quad (3.53)$$

$$= \sum_{j=1}^p \frac{\partial f_k}{\partial g_j}(\vec{g}(\vec{x})) \cdot \frac{\partial g_j}{\partial x_i}(\vec{x}) \quad (3.54)$$

$$= \sum_{j=1}^p [D\vec{f}(\vec{g}(\vec{x}))]_{k,j} \cdot [D\vec{g}(\vec{x})]_{j,i} \quad (3.55)$$

$$= \{[D\vec{f}(\vec{g}(\vec{x}))][D\vec{g}(\vec{x})]\}_{k,i}. \quad (3.56)$$

Here again we use the notation $\frac{\partial f_k}{\partial g_j}$ to denote the derivative of f_k with respect to the output of g_j . Since the output of g_j is really the j^{th} input of f_k , this derivative is just the derivative of f_k with respect to its j^{th} input.

Example 62 (Neural Networks and Backpropagation). In this example, we will develop a basic understanding of what deep neural networks are and how to train them via gradient-based optimization methods, such as those we will study in this class.

The most basic class of neural networks are “multi-layer perceptrons,” or functions of the form $\vec{f}: \mathbb{R}^n \times \Theta \rightarrow \mathbb{R}^p$, where Θ is the so-called “parameter space”, and have the form

$$\vec{f}(\vec{x}, \theta) = W^{(m)} \vec{\sigma}^{(m)}(W^{(m-1)}(\dots \vec{\sigma}^{(1)}(W^{(0)}\vec{x} + \vec{b}^{(0)}) \dots) + \vec{b}^{(m-1)}) + \vec{b}^{(m)} \quad (3.57)$$

$$\text{where } \theta = (W^{(0)}, \vec{b}^{(0)}, W^{(1)}, \vec{b}^{(1)}, W^{(2)}, \vec{b}^{(2)}, \dots, W^{(m)}, \vec{b}^{(m)}) \in \Theta, \quad (3.58)$$

and $\vec{\sigma}^{(1)}, \dots, \vec{\sigma}^{(m)}$ are “activation functions,” which we treat as generic differentiable nonlinear functions. Here θ is bolded because it should not be thought of as a scalar, vector, or matrix, but rather as a collection of matrices and vectors, an object we haven’t discussed so far in this class.

More precisely, if we define the functions $(\vec{z}^{(i)})_{i=0}^m$ and $(\vec{f}^{(i)})_{i=0}^m$ as

$$\vec{z}^{(0)}(\vec{x}, \theta) = \vec{f}^{(0)}(\vec{x}, \theta) = W^{(0)}\vec{x} + \vec{b}^{(0)} \quad (3.59)$$

$$\vec{z}^{(i)}(\vec{x}, \boldsymbol{\theta}) = \vec{f}^{(i)}(\vec{z}^{(i-1)}(\vec{x}, \boldsymbol{\theta}), \boldsymbol{\theta}) = W^{(i)}\vec{\sigma}^{(i)}(\vec{z}^{(i-1)}) + \vec{b}^{(i)}, \quad \forall i \in \{1, \dots, m\}, \quad (3.60)$$

then this formulation gives

$$\vec{f}(\vec{x}, \boldsymbol{\theta}) = \vec{z}^{(m)}(\vec{x}, \boldsymbol{\theta}) = \vec{f}^{(m)} \circ \dots \circ \vec{f}^{(0)}(\vec{x}, \boldsymbol{\theta}). \quad (3.61)$$

When we say that we “train” this neural network, we mean just that we optimize $\boldsymbol{\theta}$ to minimize some loss function evaluated on the data. Given a (differentiable) loss function $\ell: \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ and a data point $(\vec{x}, \vec{y}) \in \mathbb{R}^n \times \mathbb{R}^p$, we evaluate the loss on this data point as $\ell(\vec{f}(\vec{x}, \boldsymbol{\theta}), \vec{y})$. The true “empirical loss” across a batch of q data points (\vec{x}_i, \vec{y}_i) is

$$L(\boldsymbol{\theta}) \doteq \sum_{i=1}^q \ell(\vec{f}(\vec{x}_i, \boldsymbol{\theta}), \vec{y}_i), \quad (3.62)$$

and a “training algorithm” attempts to optimize $\boldsymbol{\theta}$ to minimize L .

For now, we concentrate on the case where we have $q = 1$, i.e., we have a single data point (\vec{x}, \vec{y}) . Here we have

$$L(\boldsymbol{\theta}) = \ell(\vec{f}(\vec{x}, \boldsymbol{\theta}), \vec{y}). \quad (3.63)$$

We now explore how to compute $\nabla L(\boldsymbol{\theta})$. In this situation, we are only differentiating with respect to $\boldsymbol{\theta}$, so we fix \vec{x} and \vec{y} , which gives the following computational graph:¹




Figure 3.5: Computational graph corresponding to the multi-layer perceptron.

Since we have not defined gradients with respect to matrices (yet), we will focus on computing $\nabla_{\vec{b}^{(i)}} L(\boldsymbol{\theta})$, i.e., the gradient of the loss function L with respect to $\vec{b}^{(i)}$, as well as $\nabla_{\vec{z}^{(i)}} L(\boldsymbol{\theta})$, i.e., the gradient of the loss function L with respect to $\vec{z}^{(i)}(\vec{x}, \boldsymbol{\theta})$. They may be interpreted as sub-vectors (i.e., literally subsets of the entries) of the full gradient $\nabla L(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta})$. This subscript notation is common in more involved problems which demand derivatives with respect to multiple vector-valued variables.

We first compute $\nabla_{\vec{z}^{(m)}} L(\boldsymbol{\theta})$. We have

$$\nabla_{\vec{z}^{(m)}} L(\boldsymbol{\theta}) = \nabla_{\vec{z}^{(m)}} \ell(\vec{z}^{(m)}, \vec{y}). \quad (3.64)$$

¹Note that we did not break the graph down to matrix multiplications, vector additions, and applications of $\vec{\sigma}^{(i)}$; we do this for clarity, because the full graph would be rather messy.

Indeed we cannot simplify this further, it being just the gradient of ℓ in its first argument.

Now we handle general $i \in \{1, \dots, m\}$. By chain rule we have the recursion:

$$\nabla_{\vec{z}^{(i-1)}} L(\boldsymbol{\theta}) = [D_{\vec{z}^{(i-1)}} L(\boldsymbol{\theta})]^\top \quad (3.65)$$

$$= [\{D_{\vec{z}^{(i)}} L(\boldsymbol{\theta})\} \{D_{\vec{z}^{(i-1)}} \vec{z}^{(i)}(\vec{x}, \boldsymbol{\theta})\}]^\top \quad (3.66)$$

$$= [D_{\vec{z}^{(i-1)}} \vec{z}^{(i)}(\vec{x}, \boldsymbol{\theta})]^\top [\nabla_{\vec{z}^{(i)}} L(\boldsymbol{\theta})]. \quad (3.67)$$

Now we want to compute $D_{\vec{z}^{(i-1)}} \vec{z}^{(i)}$. Recall that

$$\vec{z}^{(i)} = W^{(i)} \vec{\sigma}^{(i)}(\vec{z}^{(i-1)}) + \vec{b}^{(i)}. \quad (3.68)$$

For convenience, write

$$\vec{u}^{(i)} \doteq \vec{\sigma}^{(i)}(\vec{z}^{(i-1)}), \quad \text{so that} \quad \vec{z}^{(i)} = W^{(i)} \vec{u}^{(i)} + \vec{b}^{(i)}. \quad (3.69)$$

Now we compute derivatives, obtaining by the chain rule:

$$D_{\vec{z}^{(i-1)}} \vec{z}^{(i)} = [D_{\vec{u}^{(i)}} \vec{z}^{(i)}] [D_{\vec{z}^{(i-1)}} \vec{u}^{(i)}] \quad (3.70)$$

$$= W^{(i)} \cdot D\vec{\sigma}^{(i)}(\vec{z}^{(i-1)}). \quad (3.71)$$

This gives us

$$D_{\vec{z}^{(i-1)}} \vec{z}^{(i)} = W^{(i)} \cdot D\vec{\sigma}^{(i)}(\vec{z}^{(i-1)}), \quad (3.72)$$

and so

$$\nabla_{\vec{z}^{(i-1)}} L(\boldsymbol{\theta}) = [W^{(i)} \cdot D\vec{\sigma}^{(i)}(\vec{z}^{(i-1)})]^\top [\nabla_{\vec{z}^{(i)}} L(\boldsymbol{\theta})]. \quad (3.73)$$

To see the true value of this computation, we now compute $\nabla_{\vec{b}^{(i)}} L(\boldsymbol{\theta})$, for each $i \in \{0, \dots, m\}$. By another application of chain rule, we obtain

$$\nabla_{\vec{b}^{(i)}} L(\boldsymbol{\theta}) = [D_{\vec{b}^{(i)}} L(\boldsymbol{\theta})]^\top \quad (3.74)$$

$$= [\{D_{\vec{z}^{(i)}} L(\boldsymbol{\theta})\} \{D_{\vec{b}^{(i)}} \vec{z}^{(i)}\}]^\top \quad (3.75)$$

$$= [\{D_{\vec{z}^{(i)}} L(\boldsymbol{\theta})\} \cdot I]^\top \quad (3.76)$$

$$= [D_{\vec{z}^{(i)}} L(\boldsymbol{\theta})]^\top \quad (3.77)$$

$$= \nabla_{\vec{z}^{(i)}} L(\boldsymbol{\theta}). \quad (3.78)$$

Now to do automatic differentiation by backpropagation, your favorite machine learning framework (such as PyTorch) computes all the gradients $\nabla_{\vec{z}^{(i)}} L(\boldsymbol{\theta})$, starting at $i = m$ and recursing until $i = 0$. *Computing the gradients in this order saves a lot of work, since each derivative is only computed once — this is the main idea of backpropagation.* Once these derivatives are computed, one can use them to compute $\nabla_{\vec{b}^{(i)}} L(\boldsymbol{\theta})$. If we can also compute $\nabla_{W^{(i)}} L(\boldsymbol{\theta})$, then we can compute $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta})$, and thus will be able to optimize L via gradient-based optimization algorithms such as gradient descent, as discussed later in the class. However, this computation will have to wait until slightly later when we cover matrix calculus.

3.1.4 Hessian

So far, we have appropriately generalized the notion of first derivative to vector-valued functions, i.e., functions $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$. We now turn to doing the same with second derivatives.

It turns out that for general vector-valued functions \vec{f} , defining a second derivative is possible, but such an object will live in $\mathbb{R}^{m \times n \times n}$ and thus not be hard to work with using the linear algebra we have discussed in this class. However,

for multivariate functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$, defining this second derivative as a particular matrix becomes possible; this matrix, called the *Hessian*, has great conceptual and computational importance.

Recall that to find the second derivative of a scalar-valued function, we merely take the derivative of the derivative. Our notion of the gradient suffices as a first derivative; to take the derivative of this, we need to use the Jacobian. Indeed, the Hessian is exactly the Jacobian of the gradient, and defined precisely below.

Definition 63 (Hessian)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be twice differentiable. The *Hessian* of f is the function $\nabla^2 f: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ defined by

$$\nabla^2 f(\vec{x}) = D(\nabla f)(\vec{x}) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2}(\vec{x}) & \cdots & \frac{\partial^2 f}{\partial x_n \partial x_1}(\vec{x}) \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_1 \partial x_n}(\vec{x}) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(\vec{x}) \end{bmatrix}. \quad (3.79)$$

It turns out that under some mild conditions, the Hessian is symmetric; this is called Clairaut's theorem.

Theorem 64 (Clairaut's Theorem)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be twice continuously differentiable, and fix $\vec{x} \in \mathbb{R}^n$. Then $\nabla^2 f(\vec{x})$ is a symmetric matrix, i.e., for every $1 \leq i, j \leq n$ we have

$$\frac{\partial^2 f}{\partial x_i \partial x_j}(\vec{x}) = \frac{\partial^2 f}{\partial x_j \partial x_i}(\vec{x}). \quad (3.80)$$

The vast majority of functions we work with in this course are twice continuously differentiable, with some notable exceptions (i.e., the ℓ^1 norm is not even once-differentiable). Thus, in most cases, the Hessian is symmetric.

Example 65 (Hessian of Squared ℓ^2 Norm). In this example we will compute the Hessian of the function $f(\vec{x}) = \|\vec{x}\|_2^2$ where $\vec{x} \in \mathbb{R}^2$. Recall the gradient of this function computed in Example 55 as

$$\nabla f(\vec{x}) = \begin{bmatrix} 2x_1 \\ 2x_2 \end{bmatrix}. \quad (3.81)$$

The Hessian can then be computed as

$$\nabla^2 f(\vec{x}) = D(\nabla f)(\vec{x}) = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}. \quad (3.82)$$

Example 66. In this example we will compute the gradient and Hessian of the function $h(\vec{x}) = \log(1 + \|\vec{x}\|_2^2)$. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be defined by $f(x) = \log(1 + x)$, and $g: \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as $g(\vec{x}) = \|\vec{x}\|_2^2$. Then we have $Df = f'$ is the derivative of f , and $\nabla g(\vec{x}) = 2\vec{x}$ by previous examples. By the Jacobian chain rule, we have

$$\nabla h(\vec{x}) = (Dh(\vec{x}))^\top \quad (3.83)$$

$$= (D(f \circ g)(\vec{x}))^\top \quad (3.84)$$

$$= ([Df(g(\vec{x}))][Dg(\vec{x})])^\top \quad (3.85)$$

$$= ([Dg(\vec{x})]^\top [Df(g(\vec{x}))])^\top \quad (3.86)$$

$$= [\nabla g(\vec{x})][Df(g(\vec{x}))]^\top \quad (3.87)$$

$$= [\nabla g(\vec{x})][f'(g(\vec{x}))] \quad (3.88)$$

$$= \frac{2\vec{x}}{1 + \|\vec{x}\|_2^2}. \quad (3.89)$$

For the Hessian, we have

$$\nabla^2 h(\vec{x}) = D(\nabla h)(\vec{x}) \quad (3.90)$$

$$= D\left(\frac{2\vec{x}}{1 + \|\vec{x}\|_2^2}\right). \quad (3.91)$$

We compute this Jacobian, hence the desired Hessian, componentwise, and obtain

$$[\nabla^2 h(\vec{x})]_{j,k} = \left[D\left(\frac{2\vec{x}}{1 + \|\vec{x}\|_2^2}\right) \right]_{j,k} \quad (3.92)$$

$$= \frac{\partial}{\partial x_k} \left(\frac{2x_j}{1 + \|\vec{x}\|_2^2} \right) \quad (3.93)$$

$$= 2 \frac{(1 + \|\vec{x}\|_2^2) \frac{\partial}{\partial x_k} (x_j) - x_j \frac{\partial}{\partial x_k} (1 + \|\vec{x}\|_2^2)}{(1 + \|\vec{x}\|_2^2)^2} \quad (3.94)$$

$$= 2 \frac{(1 + \|\vec{x}\|_2^2) \frac{\partial}{\partial x_k} (x_j) - x_j \frac{\partial}{\partial x_k} (\|\vec{x}\|_2^2)}{(1 + \|\vec{x}\|_2^2)^2} \quad (3.95)$$

$$= 2 \frac{(1 + \|\vec{x}\|_2^2) \frac{\partial x_j}{\partial x_k} - 2x_j x_k}{(1 + \|\vec{x}\|_2^2)^2} \quad (3.96)$$

$$= -\frac{4x_j x_k}{(1 + \|\vec{x}\|_2^2)^2} + \frac{2}{1 + \|\vec{x}\|_2^2} \frac{\partial x_j}{\partial x_k} \quad (3.97)$$

$$= -\frac{4x_j x_k}{(1 + \|\vec{x}\|_2^2)^2} + \begin{cases} \frac{2}{1 + \|\vec{x}\|_2^2}, & \text{if } j = k \\ 0, & \text{if } j \neq k. \end{cases} \quad (3.98)$$

This gives

$$[\nabla^2 h(\vec{x})]_{j,k} = -\frac{4x_j x_k}{(1 + \|\vec{x}\|_2^2)^2}, \quad \forall j \neq k \quad (3.99)$$

$$[\nabla^2 h(\vec{x})]_{jj} = \frac{2}{1 + \|\vec{x}\|_2^2} - \frac{4x_j x_k}{(1 + \|\vec{x}\|_2^2)^2}, \quad \forall j. \quad (3.100)$$

We can write this using vectors as

$$\nabla^2 h(\vec{x}) = \frac{2}{1 + \|\vec{x}\|_2^2} I - \frac{4\vec{x}\vec{x}^\top}{(1 + \|\vec{x}\|_2^2)^2}. \quad (3.101)$$

3.2 Taylor's Theorems

In this section, we will introduce Taylor approximation and Taylor's theorem for the familiar scalar function case. Then we will generalize the idea of Taylor approximation to multivariate functions. Taylor approximation is a tool to find polynomial approximation of functions using information about the function value at a point along with the value of its first, second and higher order derivatives.

Definition 67 (Taylor Approximation)

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a k -times continuously differentiable function, and fix $x_0 \in \mathbb{R}$. The k^{th} degree Taylor approxi-

mation around x_0 is the function $\widehat{f}_k(\cdot; x_0): \mathbb{R} \rightarrow \mathbb{R}$ given by

$$\widehat{f}_k(x; x_0) = f(x_0) + \frac{1}{1!} \frac{df}{dx}(x_0) \cdot (x - x_0) + \cdots + \frac{1}{k!} \frac{d^k f}{dx^k}(x_0) \cdot (x - x_0)^k \quad (3.102)$$

$$= \sum_{i=0}^k \frac{1}{i!} \frac{d^i f}{dx^i}(x_0) \cdot (x - x_0)^i. \quad (3.103)$$

In particular, the first-order and second-order Taylor approximations of f around x_0 are

$$\widehat{f}_1(x; x_0) = f(x_0) + \frac{df}{dx}(x_0) \cdot (x - x_0) \quad (3.104)$$

$$\widehat{f}_2(x; x_0) = f(x_0) + \frac{df}{dx}(x_0) \cdot (x - x_0) + \frac{1}{2} \frac{d^2 f}{dx^2}(x_0) \cdot (x - x_0)^2. \quad (3.105)$$

We will derive multivariable versions of these approximations later.

Example 68 (Taylor Approximation of Cubic Function). Let us approximate the function $f(x) = x^3$ around the fixed point $x_0 = 1$ using Taylor approximations of different degrees.

$$\widehat{f}_1(x; 1) = f(x_0) + \frac{df}{dx}(x_0) \cdot (x - x_0) \quad (3.106)$$

$$= x_0^3 + 3x_0^2 \cdot (x - x_0) \quad (3.107)$$

$$= 1^3 + 3 \cdot 1^2 \cdot (x - 1) \quad (3.108)$$

$$= 3(x - 1) + 1 \quad (3.109)$$

$$= 3x - 2. \quad (3.110)$$

$$\widehat{f}_2(x; 1) = \widehat{f}_1(x; 1) + \frac{1}{2} \frac{d^2 f}{dx^2}(x_0) \cdot (x - x_0)^2 \quad (3.111)$$

$$= 3x - 2 + 3 \cdot 1 \cdot (x - 1)^2 \quad (3.112)$$

$$= 3x^2 - 3x + 1. \quad (3.113)$$

$$\widehat{f}_3(x; 1) = \widehat{f}_2(x; 1) + \frac{1}{6} \frac{d^3 f}{dx^3}(x_0) \cdot (x - x_0)^3 \quad (3.114)$$

$$= 3x^2 - 3x + 1 + (x - 1)^3 \quad (3.115)$$

$$= x^3. \quad (3.116)$$

We notice the following take-aways:

- The first-order Taylor approximation $\widehat{f}_1(\cdot; x_0)$ is the *best linear approximation* to f around $x = x_0 = 1$. In particular, its graph is the tangent line to the graph of f around the point $(x_0, f(x_0)) = (1, 1)$, as observed in Figure 3.6.
- The second-order Taylor approximation $\widehat{f}_2(\cdot; x_0)$ is the *best quadratic approximation* to f around $x = x_0 = 1$. It is the parabola whose graph passes through the point $(x_0, f(x_0)) = (1, 1)$, as observed in Figure 3.6, and it has the same first and second derivatives as f at x_0 . Using the intuition that the second derivative models curvature, we see that the second-order Taylor approximation captures the local curvature of the graph of the function. This intuition will be helpful later when discussing convexity.
- The third-degree Taylor approximation $\widehat{f}_3(\cdot; x_0)$ is the *best cubic approximation* to f ; because f is just a cubic function, the best cubic approximation is just f itself, and indeed we have $\widehat{f}_3(\cdot; x_0) = f$.




Figure 3.6: First and second degree Taylor approximations of the function $f(x) = x^3$.

Taylor approximation gives us the degree k polynomial that approximates the function $f(x)$ around the fixed point $x = x_0$. Taylor's theorem quantifies the bounds for the error of this approximation.

Theorem 69 (Taylor's Theorem)

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a function which is k -times continuously differentiable, and fix $x_0 \in \mathbb{R}$. Then for all $x \in \mathbb{R}$ we have

$$f(x) = \widehat{f}_k(x; x_0) + o(|x - x_0|^k) \quad (3.117)$$

where the term $o(|x - x_0|^k)$ (i.e., the *remainder*) denotes a function, say $R_k(x; x_0)$, such that

$$\lim_{x \rightarrow x_0} \frac{R_k(x; x_0)}{|x - x_0|^k} = 0. \quad (3.118)$$

We use this remainder notation because we don't really care about what it is precisely, only its limiting behavior as $x \rightarrow x_0$, and the little- o notation allows us to not worry too much about the exact form of the remainder.

This theorem certifies that the Taylor approximations \widehat{f}_k are good approximations to f . Another way to write this result is generally more useful or simpler:

$$f(x + \delta) = \underbrace{f(x) + \frac{df}{dx}(x) \cdot \delta}_{=\widehat{f}_1(x+\delta;x)} + o(|\delta|) \quad (3.119)$$

$$= \underbrace{f(x) + \frac{df}{dx}(x) \cdot \delta + \frac{1}{2} \frac{d^2f}{dx^2}(x) \cdot \delta^2}_{=\widehat{f}_2(x+\delta;x)} + o(\delta^2) \quad (3.120)$$

$$= \dots \quad (3.121)$$

We will never need to quantitatively work with the remainder in this course; we will usually write $f \approx \widehat{f}_k$ and leave it at that.

3.2.1 Taylor Approximation of Multivariate Functions

Using the definitions we introduced for the gradient and Hessian of multivariate functions, we can generalize the idea of Taylor's approximation to these functions.

Definition 70 (Multivariate Taylor Approximations)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and fix $\vec{x}_0 \in \mathbb{R}^n$.

- If f is continuously differentiable, then its first-order Taylor approximation around \vec{x}_0 is the function $\widehat{f}_1(\cdot; \vec{x}_0): \mathbb{R}^n \rightarrow \mathbb{R}$ given by

$$\widehat{f}_1(\vec{x}; \vec{x}_0) = f(\vec{x}_0) + [\nabla f(\vec{x}_0)]^\top (\vec{x} - \vec{x}_0). \quad (3.122)$$

- If f is twice continuously differentiable, then its second-order Taylor approximation around \vec{x}_0 is the function $\widehat{f}_2(\cdot; \vec{x}_0): \mathbb{R}^n \rightarrow \mathbb{R}$ given by

$$\widehat{f}_2(\vec{x}; \vec{x}_0) = f(\vec{x}_0) + [\nabla f(\vec{x}_0)]^\top (\vec{x} - \vec{x}_0) + \frac{1}{2}(\vec{x} - \vec{x}_0)^\top [\nabla^2 f(\vec{x}_0)](\vec{x} - \vec{x}_0). \quad (3.123)$$

The graph of the first-order Taylor approximation is the hyperplane tangent to the graph of f at the point $(\vec{x}_0, f(\vec{x}_0))$. This hyperplane has normal vector $\nabla f(\vec{x}_0)$.

We could define higher-order Taylor approximations \widehat{f}_k , but to express them concisely would require generalizations of matrices, called *tensors*. For example, the third derivative of a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is a rank-3 tensor, i.e., an object which lives in $\mathbb{R}^{n \times n \times n}$. These are out of scope for this course, and anyways we will only need the first two derivatives.

We can also state an analogous Taylor's theorem.

Theorem 71 (Taylor's Theorem)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function which is k -times continuously differentiable, and fix $\vec{x}_0 \in \mathbb{R}^n$. Then for all $\vec{x} \in \mathbb{R}^n$ we have

$$f(\vec{x}) = \widehat{f}_k(\vec{x}; \vec{x}_0) + o(\|\vec{x} - \vec{x}_0\|_2^k). \quad (3.124)$$

We can re-write this result in the following, more useful, way for $k = 1$ and $k = 2$:

$$f(\vec{x} + \vec{\delta}) = \underbrace{f(\vec{x}) + [\nabla f(\vec{x})]^\top \vec{\delta}}_{\widehat{f}_1(\vec{x} + \vec{\delta}; \vec{x})} + o(\|\vec{\delta}\|_2) \quad (3.125)$$

$$= \underbrace{f(\vec{x}) + [\nabla f(\vec{x})]^\top \vec{\delta} + \frac{1}{2} \vec{\delta}^\top [\nabla^2 f(\vec{x})] \vec{\delta}}_{\widehat{f}_2(\vec{x} + \vec{\delta}; \vec{x})} + o(\|\vec{\delta}\|_2^2) \quad (3.126)$$

$$= \dots \quad (3.127)$$

Example 72 (Taylor Approximation of the Squared ℓ^2 norm). In this example we will compute and visualize the first and second degree Taylor approximations of the squared ℓ^2 norm function $f(\vec{x}) = \|\vec{x}\|_2^2$ for $\vec{x} \in \mathbb{R}^2$ around the vector $\vec{x} = \vec{x}_0$. First recall the gradient and hessian of the function which are computed in Examples 55 and 65, respectively.

- First degree approximation:

$$\widehat{f}_1(\vec{x}; \vec{x}_0) = f(\vec{x}_0) + [\nabla f(\vec{x}_0)]^\top (\vec{x} - \vec{x}_0) \quad (3.128)$$

$$= \|\vec{x}_0\|_2^2 + [2\vec{x}_0]^\top (\vec{x} - \vec{x}_0) \quad (3.129)$$

$$= 2\vec{x}_0^\top \vec{x} - \|\vec{x}_0\|_2^2. \quad (3.130)$$

Recall that the graph of $f(\vec{x})$ has a paraboloid shape. Let us now evaluate and visualize the first degree approximation of the function around $\vec{x}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$.

$$\hat{f}_1(\vec{x}; \vec{x}_0) = 2x_1 - 1. \quad (3.131)$$

We plot this function in Figure 3.7 and notice that the graph of the first order approximation is the plane tangent to the paraboloid at the point $(1, 0, f(1, 0)) = (1, 0, 1)$.

- Second degree approximation:

$$\hat{f}_2(\vec{x}; \vec{x}_0) = \hat{f}_1(\vec{x}; \vec{x}_0) + \frac{1}{2}(\vec{x} - \vec{x}_0)^\top [\nabla^2 f(\vec{x}_0)](\vec{x} - \vec{x}_0) \quad (3.132)$$

$$= \underbrace{2\vec{x}_0^\top \vec{x} - \|\vec{x}_0\|_2^2}_{=\hat{f}_1(\vec{x}; \vec{x}_0)} + \frac{1}{2}(\vec{x} - \vec{x}_0)^\top [2I](\vec{x} - \vec{x}_0) \quad (3.133)$$

$$= 2\vec{x}_0^\top \vec{x} - \vec{x}_0^\top \vec{x}_0 + (\vec{x} - \vec{x}_0)^\top (\vec{x} - \vec{x}_0) \quad (3.134)$$

$$= 2\vec{x}_0^\top \vec{x} - \vec{x}_0^\top \vec{x}_0 + \vec{x}^\top \vec{x} - \vec{x}_0^\top \vec{x} - \vec{x}^\top \vec{x}_0 + \vec{x}_0^\top \vec{x}_0 \quad (3.135)$$

$$= 2\vec{x}_0^\top \vec{x} - \vec{x}_0^\top \vec{x}_0 + \vec{x}^\top \vec{x} - 2\vec{x}_0^\top \vec{x} + \vec{x}_0^\top \vec{x}_0 \quad (3.136)$$

$$= \vec{x}^\top \vec{x} \quad (3.137)$$

$$= \|\vec{x}\|_2^2. \quad (3.138)$$

Thus $\hat{f}_2 = f$ independently of the choice of \vec{x}_0 , which makes sense since f is a quadratic function.




Figure 3.7: First degree Taylor approximation of the function $f(\vec{x}) = \|\vec{x}\|_2^2$

Example 73. We can also compute gradients using Taylor's theorem by pattern matching; this is sometimes much neater than taking componentwise gradients. At first glance this seems circular, but we will see how it is possible. Take for example the function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ given by $f(\vec{x}) = \vec{x}^\top A \vec{x}$. We can perturb f around \vec{x} to obtain

$$f(\vec{x} + \vec{\delta}) = (\vec{x} + \vec{\delta})^\top A(\vec{x} + \vec{\delta}) \quad (3.139)$$

$$= \vec{x}^\top A \vec{x} + \vec{\delta}^\top A \vec{x} + \vec{x}^\top A \vec{\delta} + \vec{\delta}^\top A \vec{\delta} \quad (3.140)$$

$$= f(\vec{x}) + (\vec{x}^\top A^\top + \vec{x}^\top A) \vec{\delta} + \vec{\delta}^\top A \vec{\delta} \quad (3.141)$$

$$= f(\vec{x}) + ((A + A^\top) \vec{x})^\top \vec{\delta} + \frac{1}{2} \vec{\delta}^\top (A + A^\top) \vec{\delta}. \quad (3.142)$$

However, Taylor's theorem tells us that

$$f(\vec{x} + \vec{\delta}) = f(\vec{x}) + [\nabla f(\vec{x})]^\top \vec{\delta} + \frac{1}{2} \vec{\delta}^\top [\nabla^2 f(\vec{x})] \vec{\delta} + o(\|\vec{\delta}\|_2^2). \quad (3.143)$$

By pattern matching, we see that $\nabla f(\vec{x}) = (A + A^\top)\vec{x}$ (as obtained in a previous example), $\nabla^2 f(\vec{x}) = A + A^\top$, and the remainder term is $\vec{0}$.

One final note is that we changed $2A \rightarrow A + A^\top$ in Equation (3.142); this is to ensure that the Hessian is symmetric, which is a consequence of Theorem 64; and we are able to do this because

$$\vec{\delta}^\top (2A)\vec{\delta} = \vec{\delta}^\top A\vec{\delta} + \vec{\delta}^\top A\vec{\delta} = \vec{\delta}^\top A\vec{\delta} + (\vec{\delta}^\top A\vec{\delta})^\top = \vec{\delta}^\top A\vec{\delta} + \vec{\delta}^\top A^\top \vec{\delta} = \vec{\delta}^\top (A + A^\top)\vec{\delta}. \quad (3.144)$$

We conclude by introducing a more general version of a first-order Taylor approximation, a corresponding Taylor's theorem, and giving an example of when it is useful.

Definition 74 (Vector-Valued Taylor Approximation)

Let $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and fix $\vec{x}_0 \in \mathbb{R}^n$. If \vec{f} is continuously differentiable, then its first-order Taylor approximation around \vec{x}_0 is the function $\hat{f}_1: \mathbb{R}^n \rightarrow \mathbb{R}^m$ given by

$$\hat{f}_1(\vec{x}; \vec{x}_0) = \vec{f}(\vec{x}_0) + [D\vec{f}(\vec{x}_0)](\vec{x} - \vec{x}_0). \quad (3.145)$$

Again, higher-order approximations will require higher-order derivatives, which requires tensors.

Theorem 75 (Vector-Valued Taylor's Theorem)

Let $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a continuously differentiable function, and fix $\vec{x}_0 \in \mathbb{R}^n$. Then for all $\vec{x} \in \mathbb{R}^n$ we have

$$\vec{f}(\vec{x}) = \hat{f}_1(\vec{x}; \vec{x}_0) + o(\|\vec{x} - \vec{x}_0\|_2). \quad (3.146)$$

We can again re-write this result in a more workable form:

$$\vec{f}(\vec{x} + \vec{\delta}) = \vec{f}(\vec{x}) + [D\vec{f}(\vec{x})]\vec{\delta} + o(\|\vec{\delta}\|_2). \quad (3.147)$$

Example 76. Taylor's theorem can be used to compute gradients by pattern matching, even when the function is not linear or quadratic. For instance, we now use it to derive the chain rule (albeit with stronger assumptions on the functions). Let $\vec{f}: \mathbb{R}^p \rightarrow \mathbb{R}^m$ and $\vec{g}: \mathbb{R}^n \rightarrow \mathbb{R}^p$ be continuously differentiable. Let $\vec{h}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be defined as $\vec{h}(\vec{x}) = \vec{f}(\vec{g}(\vec{x}))$. Then we compute \vec{h} on a perturbation around \vec{x} and expand:

$$\vec{h}(\vec{x} + \vec{\delta}) = \vec{f}(\vec{g}(\vec{x} + \vec{\delta})) \quad (3.148)$$

$$\approx \vec{f}(\vec{g}(\vec{x}) + [D\vec{g}(\vec{x})]\vec{\delta}) \quad (3.149)$$

$$\approx \vec{f}(\vec{g}(\vec{x})) + [D\vec{f}(\vec{g}(\vec{x}))]([D\vec{g}(\vec{x})]\vec{\delta}) \quad (3.150)$$

$$\approx \vec{f}(\vec{g}(\vec{x})) + [D\vec{f}(\vec{g}(\vec{x}))][D\vec{g}(\vec{x})]\vec{\delta}. \quad (3.151)$$

The first Taylor expansion is an expansion of \vec{g} around the point \vec{x} with perturbation $\vec{\delta}$; the second Taylor expansion is an expansion of \vec{f} around the point $\vec{g}(\vec{x})$ with perturbation $[D\vec{g}(\vec{x})]\vec{\delta}$.

Meanwhile, Taylor's theorem says that

$$\vec{h}(\vec{x} + \vec{\delta}) \approx \vec{h}(\vec{x}) + [D\vec{h}(\vec{x})]\vec{\delta}. \quad (3.152)$$

Thus by pattern matching we find that

$$D\vec{h}(\vec{x}) = [D\vec{f}(\vec{g}(\vec{x}))][D\vec{g}(\vec{x})] \quad (3.153)$$

which is precisely the chain rule!

Note that here we did not invoke the little- o notation because it turns out to be quite messy, but it is indeed possible to do the required rigorous manipulations and get the same result.

As a last practical note, remembering the formula for Taylor approximations helps us confirm our understanding of the dimensions of each vector. For instance, every term should multiply to a scalar. This makes it simpler to remember that, for a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, the gradient ∇f outputs column vectors in \mathbb{R}^n , the Hessian $\nabla^2 f$ outputs square matrices in $\mathbb{R}^{n \times n}$, etc.

3.3 The Main Theorem

In this section, we will use the concepts introduced in the previous sections to state and prove one of the fundamental ideas in optimization.

Theorem 77 (The Main Theorem [4])

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function, and let $\Omega \subseteq \mathbb{R}^n$ be an open set.^a Consider the optimization problem

$$\min_{\vec{x} \in \Omega} f(\vec{x}). \quad (3.154)$$

Let \vec{x}^* be a solution to this optimization problem. Then

$$\nabla f(\vec{x}^*) = \vec{0}. \quad (3.155)$$

^a“Open sets” are analogous to open intervals (a, b) — i.e., not containing boundary points.

This theorem gives a *necessary* condition for a point to be an optimal solution of this optimization problem. It says that any point that is optimal must necessarily have gradient equal to zero.

Proof. We prove this for scalar functions $f: \mathbb{R} \rightarrow \mathbb{R}$ only; the vector case is a bit more complicated and is left as an exercise.

Using Taylor approximation of the function around the optimal point:

$$f(x) = f(x^*) + \frac{df}{dx}(x^*) \cdot (x - x^*) + o(|x - x^*|). \quad (3.156)$$

Since $f(x^*) \leq f(x)$ for all $x \in \Omega$, we have

$$f(x) \leq f(x) + \frac{df}{dx}(x^*) \cdot (x - x^*) + o(|x - x^*|) \quad (3.157)$$

$$\implies 0 \leq \frac{df}{dx}(x^*) \cdot (x - x^*) + o(|x - x^*|). \quad (3.158)$$

Since Ω is an open set, there exists some ball of positive radius $r > 0$ around x^* such that $B_r(x^*) \subseteq \Omega$. Formally,

$$B_r(x^*) = \{x \in \mathbb{R} \mid |x - x^*| \leq r\}. \quad (3.159)$$

Let us partition $B_r(x^*)$ into B_+ , the set of all $x \in B_r(x^*)$ such that $x - x^* \geq 0$, and B_- , the set of all $x \in B_r(x^*)$ such that $x - x^* < 0$.

For all $x \in B_+$, we have

$$0 \leq \frac{df}{dx}(x^*) \cdot (x - x^*) + o(|x - x^*|) \quad (3.160)$$

$$= \frac{df}{dx}(x^*) \cdot |x - x^*| + o(|x - x^*|) \quad (3.161)$$

$$\implies 0 \leq \frac{df}{dx}(x^*) + \frac{o(|x - x^*|)}{|x - x^*|}. \quad (3.162)$$

Taking the limit as $x \rightarrow x^*$ within B_+ , we have

$$0 \leq \lim_{\substack{x \rightarrow x^* \\ x \in B_+}} \left\{ \frac{df}{dx}(x^*) + \frac{o(|x - x^*|)}{|x - x^*|} \right\} \quad (3.163)$$

$$= \frac{df}{dx}(x^*) + \lim_{\substack{x \rightarrow x^* \\ x \in B_+}} \frac{o(|x - x^*|)}{|x - x^*|} \quad (3.164)$$

$$= \frac{df}{dx}(x^*). \quad (3.165)$$

Thus we have $0 \leq \frac{df}{dx}(x^*)$. On the other hand, for all $x \in B_-$, we have

$$0 \leq \frac{df}{dx}(x^*) \cdot (x - x^*) + o(|x - x^*|) \quad (3.166)$$

$$= -\frac{df}{dx}(x^*) \cdot |x - x^*| + o(|x - x^*|) \quad (3.167)$$

$$\implies 0 \geq \frac{df}{dx}(x^*) - \frac{o(|x - x^*|)}{|x - x^*|}. \quad (3.168)$$

Taking the limit $x \rightarrow x^*$ within B_- , we have

$$0 \geq \lim_{\substack{x \rightarrow x^* \\ x \in B_-}} \left\{ \frac{df}{dx}(x^*) - \frac{o(|x - x^*|)}{|x - x^*|} \right\} \quad (3.169)$$

$$= \frac{df}{dx}(x^*) - \lim_{\substack{x \rightarrow x^* \\ x \in B_-}} \frac{o(|x - x^*|)}{|x - x^*|} \quad (3.170)$$

$$= \frac{df}{dx}(x^*). \quad (3.171)$$

Thus we have $0 \geq \frac{df}{dx}(x^*)$ and so $\frac{df}{dx}(x^*) = 0$. □

3.4 Directional Derivatives

Recall the definition of the partial derivative of a multivariate function (Definition 47), which represents the rate of change of the function $f(\vec{x})$ along one of the standard basis vectors. We do not need to restrict our treatment to the standard basis vectors; in fact, we can compute the rate of change of the function in any arbitrary direction. This is called the directional derivative.

Definition 78 (Directional Derivative)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable, and fix $\vec{u} \in \mathbb{R}^n$ such that $\|\vec{u}\|_2 = 1$. The *directional derivative* of f along \vec{u} is

the function $Df(\cdot)[\vec{u}]: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$Df(\vec{x})[\vec{u}] = \lim_{h \rightarrow 0} \frac{f(\vec{x} + h \cdot \vec{u}) - f(\vec{x})}{h}. \quad (3.172)$$

If we know the directional derivative in any direction, we know the gradient; similarly, if we know the gradient, we know the directional derivative. The way to connect the two is given by the following proposition, whose proof is left as an exercise.

Proposition 79

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable, and fix $\vec{u} \in \mathbb{R}^n$ such that $\|\vec{u}\|_2 = 1$. Then

$$Df(\vec{x})[\vec{u}] = \vec{u}^\top [\nabla f(\vec{x})]. \quad (3.173)$$

In particular, $Df(\vec{x})[\vec{e}_i] = \frac{\partial f}{\partial x_i}(\vec{x})$.

3.5 (OPTIONAL) Matrix Calculus

So far we have only discussed derivatives of three types of function and all have either scalars or vectors and their input and output. We can think of a more general class of functions that also involve matrices. In this section, we will generalize the idea of derivatives to such functions. We will focus our attention on functions of the form $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$, which take a matrix $X \in \mathbb{R}^{m \times n}$ as input and produce a scalar $f(X)$ as output. Familiar examples of such functions include matrix norms, the determinant, and the trace.

Definition 80 (Gradient)

Let $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be differentiable. The *gradient* of f is the function $\nabla f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ which is defined as

$$\nabla f(X) = \begin{bmatrix} \frac{\partial f}{\partial X_{11}}(X) & \cdots & \frac{\partial f}{\partial X_{1n}}(X) \\ \vdots & \ddots & \vdots \\ \frac{\partial f}{\partial X_{m1}}(X) & \cdots & \frac{\partial f}{\partial X_{mn}}(X) \end{bmatrix} \quad (3.174)$$

There exists a general chain rule for matrix-valued functions, which is provable by flattening out all matrices into vectors and applying the vector chain rule.

Theorem 81 (Chain Rule)

Let $F: \mathbb{R}^{p \times q} \rightarrow \mathbb{R}^{r \times s}$ and $G: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{p \times q}$ be differentiable functions. Let $H: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{r \times s}$ be defined by $H(X) = F(G(X))$ for all $X \in \mathbb{R}^{m \times n}$. Then H is differentiable, and for all i, j, k, ℓ , we have

$$\frac{\partial H_{ij}}{\partial X_{k\ell}}(X) = \sum_a \sum_b \frac{\partial F_{ij}}{\partial G_{ab}}(G(X)) \frac{\partial G_{ab}}{\partial X_{k\ell}}(X) \quad (3.175)$$

As before, the notation $\frac{\partial F_{ij}}{\partial G_{ab}}$ means to take the derivative of the ij^{th} output of F by its ab^{th} input. A more specific version of this chain rule is given below for functions $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$.

Proposition 82

Let $F: \mathbb{R}^{p \times q} \rightarrow \mathbb{R}$ and $G: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{p \times q}$ be differentiable functions. Let $h: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be defined by $h(X) = f(G(X))$ for all $X \in \mathbb{R}^{m \times n}$. Then h is differentiable, and for all k, ℓ , we have

$$\frac{\partial h}{\partial X_{k\ell}}(X) = \sum_a \sum_b [\nabla f(G(X))]_{ab} \frac{\partial G_{ab}}{\partial X_{k\ell}}(X) \quad (3.176)$$

We also are able to define a first-order Taylor expansion without having to use tensor notation.

Definition 83 (Matrix Taylor Approximation)

Let $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ and fix $X_0 \in \mathbb{R}^{m \times n}$. If f is continuously differentiable, then its first-order Taylor approximation around X_0 is the function $\hat{f}_1(\cdot; X_0): \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ given by

$$\hat{f}_1(X; X_0) = f(X_0) + \text{tr}([\nabla f(X_0)]^\top (X - X_0)). \quad (3.177)$$

There is a corresponding Taylor's theorem certifying the Taylor approximation accuracy, but we don't state it here.

Finally, note that the general recipe for computing all quantities such as the gradient, Jacobian, and gradient matrix is the same: consider each input component and each output component separately and organize their partial derivatives in vector or matrix form with a standard layout.

Example 84 (Finishing Example 62). Now that we know how to take matrix-valued gradients, we complete the example of [Neural Networks and Backpropagation](#). Before reading the following, please revise the lengthy setup of this example.

We promised in this example a way to compute $\nabla_{\theta} L(\theta)$, or more precisely a way to compute $\nabla_{W^{(i)}} L(\theta)$. We now have the tools to do this using the chain rule. Recall that we have access to $\nabla_{\vec{z}^{(i)}} L(\theta)$ by backpropagation. Then we can compute the components of $\nabla_{W^{(i)}} L(\theta)$ by

$$[\nabla_{W^{(i)}} L(\theta)]_{j,k} = \frac{\partial L}{\partial (W^{(i)})_{j,k}}(\theta) \quad (3.178)$$

$$= \sum_a [\nabla_{\vec{z}^{(i)}} L(\theta)]_a \frac{\partial (\vec{z}^{(i)})_a}{\partial (W^{(i)})_{j,k}} \quad (3.179)$$

$$= \sum_a [\nabla_{\vec{z}^{(i)}} L(\theta)]_a \cdot \begin{cases} [\vec{\sigma}^{(i)}(\vec{z}^{(i-1)})]_k, & \text{if } a = j \text{ and } i \in \{1, \dots, m\} \\ x_k, & \text{if } a = j \text{ and } i = 0 \\ 0, & \text{otherwise} \end{cases} \quad (3.180)$$

$$= \begin{cases} [\nabla_{\vec{z}^{(i)}} L(\theta)]_j [\vec{\sigma}^{(i)}(\vec{z}^{(i-1)})]_k, & \text{if } i \in \{1, \dots, m\} \\ [\nabla_{\vec{z}^{(i)}} L(\theta)]_j [\vec{x}]_k, & \text{if } i = 0 \end{cases} \quad (3.181)$$

$$= \begin{cases} \{[\nabla_{\vec{z}^{(i)}} L(\theta)]\} \vec{\sigma}^{(i)}(\vec{z}^{(i-1)})^\top]_{j,k}, & \text{if } i \in \{1, \dots, m\} \\ \{[\nabla_{\vec{z}^{(i)}} L(\theta)]\} \vec{x}^\top]_{j,k}, & \text{if } i = 0. \end{cases} \quad (3.182)$$

This gives

$$\nabla_{W^{(i)}} L(\theta) = \begin{cases} [\nabla_{\vec{z}^{(i)}} L(\theta)] \vec{\sigma}^{(i)}(\vec{z}^{(i-1)})^\top, & \text{if } i \in \{1, \dots, m\} \\ [\nabla_{\vec{z}^{(i)}} L(\theta)] \vec{x}^\top, & \text{if } i = 0. \end{cases} \quad (3.183)$$

In combination with the expression for $\nabla_{\vec{b}^{(i)}}$ from Example 62, we can efficiently compute $\nabla_{\theta} L(\theta)$, and are able to train our neural network via gradient-based optimization methods such as gradient descent.

Chapter 4

Linear and Ridge Regression

Relevant sections of the textbooks:

- [2] Chapter 6.

4.1 Impact of Perturbations on Linear Regression

Before we start thinking about generic convex analysis, we will first study a particularly instructive, useful, and interesting linear-algebraic convex optimization problem.

Let $A \in \mathbb{R}^{n \times n}$ be invertible and $\vec{y} \in \mathbb{R}^n$. Consider the generic linear system $A\vec{x} = \vec{y}$, perhaps representing measurements of some physical system. There is exactly one \vec{x} which solves this system — that being $\vec{x} = A^{-1}\vec{y}$. We want to understand how sensitive this system is to perturbations in the output. That is, if \vec{y} is perturbed by $\vec{\delta}_{\vec{y}}$ for $\|\vec{\delta}_{\vec{y}}\|_2$ small (say, representing noise in the measurements), then the \vec{x} that solves the system is *also* perturbed, say by $\vec{\delta}_{\vec{x}}$. So in the end, we have

$$A(\vec{x} + \vec{\delta}_{\vec{x}}) = (\vec{y} + \vec{\delta}_{\vec{y}}). \quad (4.1)$$

We want to compute the relative change in \vec{x} , that is, $\frac{\|\vec{\delta}_{\vec{x}}\|_2}{\|\vec{x}\|_2}$, in terms of $\|\vec{\delta}_{\vec{y}}\|_2$, as well as other properties of the system. In the context of our physical measurement system, we would much rather have this ratio be *small*; this means that the solutions to the equations governing our physical system are *robust* to measurement errors, thus assuring us that our model is relatively accurate to the real-life physical system. Thus, at the least we want to upper-bound $\frac{\|\vec{\delta}_{\vec{x}}\|_2}{\|\vec{x}\|_2}$.

The first part of upper-bounding $\frac{\|\vec{\delta}_{\vec{x}}\|_2}{\|\vec{x}\|_2}$ is to upper-bound $\|\vec{\delta}_{\vec{x}}\|_2$. We have

$$A(\vec{x} + \vec{\delta}_{\vec{x}}) = \vec{y} + \vec{\delta}_{\vec{y}} \quad (4.2)$$

$$A\vec{x} + A\vec{\delta}_{\vec{x}} = \vec{y} + \vec{\delta}_{\vec{y}} \quad (4.3)$$

$$A\vec{\delta}_{\vec{x}} = \vec{\delta}_{\vec{y}} \quad (4.4)$$

$$\vec{\delta}_{\vec{x}} = A^{-1}\vec{\delta}_{\vec{y}}. \quad (4.5)$$

Then by taking norms on both sides,

$$\|\vec{\delta}_{\vec{x}}\|_2 = \|A^{-1}\vec{\delta}_{\vec{y}}\|_2 \quad (4.6)$$

$$\leq \max_{\substack{\vec{z} \in \mathbb{R}^n \\ \|\vec{z}\|_2 = \|\vec{\delta}_{\vec{y}}\|_2}} \|A^{-1}\vec{z}\|_2 \quad (4.7)$$

$$= \left(\max_{\substack{\vec{z} \in \mathbb{R}^n \\ \|\vec{z}\|_2=1}} \|A^{-1}\vec{z}\|_2 \right) \|\vec{\delta}_{\vec{y}}\|_2 \quad (4.8)$$

$$= \|A^{-1}\|_2 \|\vec{\delta}_{\vec{y}}\|_2. \quad (4.9)$$

In order to upper-bound $\frac{\|\vec{\delta}_{\vec{x}}\|_2}{\|\vec{x}\|_2}$, we also need to lower-bound $\|\vec{x}\|_2$. Applying the same matrix norm inequality to the regular linear system $A\vec{x} = \vec{y}$ gives

$$A\vec{x} = \vec{y} \quad (4.10)$$

$$\|A\vec{x}\|_2 = \|\vec{y}\|_2 \quad (4.11)$$

$$\|A\|_2 \|\vec{x}\|_2 \geq \|\vec{y}\|_2 \quad (4.12)$$

$$\|\vec{x}\|_2 \geq \frac{\|\vec{y}\|_2}{\|A\|_2} \quad (4.13)$$

where $\|A\|_2 \neq 0$ because A is invertible. Plugging in both bounds, we have

$$\frac{\|\vec{\delta}_{\vec{x}}\|_2}{\|\vec{x}\|_2} \leq \frac{\|A^{-1}\|_2 \|\vec{\delta}_{\vec{y}}\|_2}{\|\vec{y}\|_2 / \|A\|_2} \quad (4.14)$$

$$= \|A\|_2 \|A^{-1}\|_2 \cdot \frac{\|\vec{\delta}_{\vec{y}}\|_2}{\|\vec{y}\|_2}. \quad (4.15)$$

Thus we've bounded the relative change in \vec{x} by the relative change in \vec{y} . If the relative change in \vec{y} is small, then the relative change in \vec{x} will be small, and so on. But we'd like to say something more about $\|A\|_2 \|A^{-1}\|_2$, and indeed we can:

$$\|A\|_2 \|A^{-1}\|_2 = \sigma_1\{A\} \cdot \sigma_1\{A^{-1}\} = \frac{\sigma_1\{A\}}{\sigma_n\{A\}}, \quad (4.16)$$

where again, $\sigma_n\{A\} \neq 0$ because A is invertible. This quantity

$$\kappa(A) \doteq \frac{\sigma_1\{A\}}{\sigma_n\{A\}} \quad (4.17)$$

is called the *condition number* of a matrix. In general, for non-invertible systems, this can be infinite, but has the same definition.

Definition 85 (Condition Number)

Let $A \in \mathbb{R}^{n \times n}$. The *condition number* of A , denoted $\kappa(A)$, is given by

$$\kappa(A) \doteq \frac{\sigma_1\{A\}}{\sigma_n\{A\}}. \quad (4.18)$$

If $\kappa(A)$ is large, then even a small change in our measurement \vec{y} will result in a huge change in our variable \vec{x} . If $\kappa(A)$ is small, then large changes in our measurement \vec{y} result in small changes to our variable \vec{x} .

It seems unlikely that in general, the equations that define our system will be square. Most likely we will have a least-squares type *tall* system. But this is resolved by using the so-called *normal equations* to represent the least squares solution:

$$A^\top A \vec{x} = A^\top \vec{y}. \quad (4.19)$$

The condition number of this linear system is $\kappa(A^\top A)$. Since $A^\top A$ is symmetric and positive semidefinite, its eigenvalues are also its singular values, and so we have

$$\kappa(A^\top A) = \frac{\lambda_{\max}\{A^\top A\}}{\lambda_{\min}\{A^\top A\}}. \quad (4.20)$$

4.2 Ridge Regression

Sometimes we have least squares systems with $\kappa(A^\top A) = \infty$, or even finite but very large. How do we make this system solvable, robust, or even better conditioned? The answer goes like the following: suppose we could add some number λ to all eigenvalues of $A^\top A$. Then, since $\lambda_1\{A^\top A\}$ and $\lambda_n\{A^\top A\}$ go up by the same amount, $\lambda_n\{A^\top A\}$ becomes a larger fraction of $\lambda_1\{A^\top A\}$, so the condition number $\kappa(A^\top A)$ becomes lower.

This is perhaps easier to see with a numerical example, which we provide now. Suppose that $A^\top A$ has $\lambda_1\{A^\top A\} = 5$ and $\lambda_n\{A^\top A\} = 0.01$. Then $\kappa(A^\top A) = 500$. But if we add 3 to all eigenvalues of $A^\top A$, then $\lambda_1\{A^\top A\} = 8$ and $\lambda_n\{A^\top A\} = 3.01$, so $\kappa(A^\top A) = \frac{8}{3.01} \approx 2.65$. This is a much better-conditioned problem.

The question is now how to add λ to all eigenvalues of $A^\top A$. Using the shift property of eigenvalues, we see that we can add λI to $A^\top A$, so that instead of solving the system $A^\top A \vec{x} = A^\top \vec{y}$ we instead solve the system $(A^\top A + \lambda I) \vec{x} = A^\top \vec{y}$. The problem of finding \vec{x} which solves this system is called *ridge regression*. It turns out to be equivalent to the following formulation.

Theorem 86 (Ridge Regression)

Let $A \in \mathbb{R}^{m \times n}$, $\vec{y} \in \mathbb{R}^m$, and $\lambda > 0$. The unique solution to the *ridge regression* problem

$$\min_{\vec{x} \in \mathbb{R}^n} \left\{ \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_2^2 \right\} \quad (4.21)$$

is given by

$$\vec{x}^* = (A^\top A + \lambda I)^{-1} A^\top \vec{y}. \quad (4.22)$$

Proof. Let $f(\vec{x}) \doteq \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_2^2$. By taking gradients, we get

$$\nabla_{\vec{x}} f(\vec{x}) = \nabla_{\vec{x}} \left\{ \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_2^2 \right\} \quad (4.23)$$

$$= \nabla_{\vec{x}} \{ \vec{x}^\top A^\top A \vec{x} - 2\vec{y}^\top A \vec{x} + \vec{y}^\top \vec{y} + \lambda \vec{x}^\top \vec{x} \} \quad (4.24)$$

$$= 2A^\top A \vec{x} - 2A^\top \vec{y} + 2\lambda \vec{x} \quad (4.25)$$

$$= 2(A^\top A + \lambda I) \vec{x} - 2A^\top \vec{y}. \quad (4.26)$$

Thus we get that the optimal point is determined by solving the linear system

$$(A^\top A + \lambda I) \vec{x} = A^\top \vec{y}. \quad (4.27)$$

Since $A^\top A$ is PSD and $\lambda > 0$, we have $A^\top A + \lambda I$ is PD and thus invertible. Therefore

$$\vec{x}^* = (A^\top A + \lambda I)^{-1} A^\top \vec{y} \quad (4.28)$$

is the unique solution to the above linear system and therefore the unique solution to the optimization problem. \square

Note that we haven't proved that a convex function (such as the above ridge regression objective) is minimized when its derivative is 0; we prove this in subsequent lectures, but for now let us take it for granted.

Proof. Another way to solve the same problem is to consider the augmented system

$$\begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \vec{x} = \begin{bmatrix} \vec{y} \\ \vec{0} \end{bmatrix}. \quad (4.29)$$

This augmented matrix has full column rank, so we can use the least squares solution to get a unique solution for \vec{x} . We get

$$\vec{x} = \left(\begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix}^\top \begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \right)^{-1} \begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix}^\top \begin{bmatrix} \vec{y} \\ \vec{0} \end{bmatrix} \quad (4.30)$$

$$= \left(\begin{bmatrix} A^\top & \sqrt{\lambda}I \end{bmatrix} \begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \right)^{-1} \begin{bmatrix} A^\top & \sqrt{\lambda}I \end{bmatrix} \begin{bmatrix} \vec{y} \\ \vec{0} \end{bmatrix} \quad (4.31)$$

$$= \left(\begin{bmatrix} A^\top & \sqrt{\lambda}I \end{bmatrix} \begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \right)^{-1} (A^\top \vec{y} + \sqrt{\lambda}I \cdot \vec{0}) \quad (4.32)$$

$$= \left(\begin{bmatrix} A^\top & \sqrt{\lambda}I \end{bmatrix} \begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \right)^{-1} A^\top \vec{y} \quad (4.33)$$

$$= (A^\top A + \lambda I)^{-1} A^\top \vec{y}. \quad (4.34)$$

□

In the ridge regression objective

$$\|A\vec{x} - \vec{b}\|_2^2 + \lambda \|\vec{x}\|_2^2, \quad (4.35)$$

the second term $\lambda \|\vec{x}\|_2^2$ is called a *regularizer*; this is because it *regularizes* or *regularizes* our problem by making it better-conditioned.

4.3 Principal Components Regression

We can gain more understanding of the ridge regression solution by looking at it through the SVD of A . Indeed, let $A = U\Sigma V^\top$. Then the ridge regression solution is

$$x^* = (A^\top A + \lambda I)^{-1} A^\top \vec{y} \quad (4.36)$$

$$= ((U\Sigma V^\top)^\top (U\Sigma V^\top) + \lambda I)^{-1} (U\Sigma V^\top)^\top \vec{y} \quad (4.37)$$

$$= (V\Sigma^\top U^\top U\Sigma V^\top + \lambda I)^{-1} V\Sigma^\top U^\top \vec{y} \quad (4.38)$$

$$= (V\Sigma^\top \Sigma V^\top + \lambda I)^{-1} V\Sigma^\top U^\top \vec{y} \quad (4.39)$$

$$= (V(\Sigma^\top \Sigma + \lambda I) V^\top)^{-1} V\Sigma^\top U^\top \vec{y} \quad (4.40)$$

$$= (V(\Sigma^\top \Sigma + \lambda I) V^\top)^{-1} V\Sigma^\top U^\top \vec{y} \quad (4.41)$$

$$= V(\Sigma^\top \Sigma + \lambda I)^{-1} V^\top V\Sigma^\top U^\top \vec{y} \quad (4.42)$$

$$= V(\Sigma^\top \Sigma + \lambda I)^{-1} \Sigma^\top U^\top \vec{y} \quad (4.43)$$

$$= V \begin{bmatrix} (\Sigma_r^2 + \lambda I)^{-1} \Sigma_r & 0 \\ 0 & 0 \end{bmatrix} U^\top \vec{y}. \quad (4.44)$$

Looking at the middle matrix a bit more, we see that

$$(\Sigma_r^2 + \lambda I)^{-1} \Sigma_r = \begin{bmatrix} \frac{\sigma_1\{A\}}{\sigma_1\{A\}^2 + \lambda} & & \\ & \ddots & \\ & & \frac{\sigma_n\{A\}}{\sigma_n\{A\}^2 + \lambda} \end{bmatrix}.$$

Thus, we get

$$\vec{x}^* = V \begin{bmatrix} (\Sigma_r^2 + \lambda I)^{-1} \Sigma_r & 0 \\ 0 & 0 \end{bmatrix} U^\top \vec{y} \quad (4.45)$$

$$= \left(\sum_{i=1}^r \frac{\sigma_i\{A\}}{\sigma_i\{A\}^2 + \lambda} \vec{v}_i \vec{u}_i^\top \right) \vec{y} \quad (4.46)$$

$$= \sum_{i=1}^r \frac{\sigma_i\{A\}}{\sigma_i\{A\}^2 + \lambda} (\vec{u}_i^\top \vec{y}) \cdot \vec{v}_i. \quad (4.47)$$

To understand what λ is doing here, we contrast two examples. Let $A \in \mathbb{R}^{n \times 3}$ for some large $n \gg 3$.

Suppose first that $\sigma_1\{A\} = \sigma_2\{A\} = \sigma_3\{A\} = 1$. Then

$$\vec{x}^* = \frac{\sigma_1\{A\}}{\sigma_1\{A\}^2 + \lambda} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{\sigma_2\{A\}}{\sigma_2\{A\}^2 + \lambda} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{\sigma_3\{A\}}{\sigma_3\{A\}^2 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3 \quad (4.48)$$

$$= \frac{1}{1 + \lambda} \{ (\vec{u}_1^\top \vec{y}) \vec{v}_1 + (\vec{u}_2^\top \vec{y}) \vec{v}_2 + (\vec{u}_3^\top \vec{y}) \vec{v}_3 \} \quad (4.49)$$

$$= \frac{1}{1 + \lambda} \vec{\tilde{x}} \quad (4.50)$$

where $\vec{\tilde{x}}$ is the solution of the corresponding least squares linear regression problem, namely the ridge problem with $\lambda = 0$. In this way, the λ parameter decays the solution in each principal direction equally, pulling the whole $\vec{\tilde{x}}$ vector towards 0. This is interesting precisely because a first-level examination of the ridge regression objective function — and namely the $\|\vec{x}\|_2^2$ term, which by itself penalizes every direction of \vec{x} equally — may make it seem like this is always the case, but it turns out to not be, as we will see shortly.

Now suppose that $\sigma_1\{A\} = 100$, $\sigma_2\{A\} = 10$, and $\sigma_3\{A\} = 1$. Then

$$\vec{x}^* = \frac{\sigma_1\{A\}}{\sigma_1\{A\}^2 + \lambda} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{\sigma_2\{A\}}{\sigma_2\{A\}^2 + \lambda} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{\sigma_3\{A\}}{\sigma_3\{A\}^2 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3 \quad (4.51)$$

$$= \frac{100}{10000 + \lambda} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{10}{100 + \lambda} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{1}{1 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3 \quad (4.52)$$

$$= \frac{1}{100 + \lambda/100} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{1}{10 + \lambda/10} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{1}{1 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3. \quad (4.53)$$

Thus, the different terms are now impacted differently based on λ ; in particular, to impact the first term by a certain amount, one needs to change λ by 100 times the amount required to change the last term. Namely, if we set λ to be large, say $\lambda = 10000$, the coefficient of the first term becomes $1/110$, while the coefficient of the last term becomes $1/10001$ which is much lower. More generally, for a larger example, setting λ to be large effectively zeros out the last few terms while effectively not changing the first few terms. Thus, setting λ to be large effectively performs a “soft thresholding” of the singular values, making the terms associated with smaller singular values be nearly 0 while preserving the terms associated with larger singular values. More quantitatively, for large λ , we have

$$\vec{x}^* = \frac{1}{100 + \lambda/100} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{1}{10 + \lambda/10} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{1}{1 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3 \quad (4.54)$$

$$\approx \frac{1}{100 + \lambda/100} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{1}{10 + \lambda/10} (\vec{u}_2^\top \vec{y}) \vec{v}_2 \quad (4.55)$$

and for even larger λ we simply have

$$\vec{x}^* = \frac{1}{100 + \lambda/100} (\vec{u}_1^\top \vec{y}) \vec{v}_1 + \frac{1}{10 + \lambda/10} (\vec{u}_2^\top \vec{y}) \vec{v}_2 + \frac{1}{1 + \lambda} (\vec{u}_3^\top \vec{y}) \vec{v}_3 \quad (4.56)$$

$$\approx \frac{1}{100 + \lambda/100} (\vec{u}_1^\top \vec{y}) \vec{v}_1. \quad (4.57)$$

Since the terms form a linear combination of the \vec{v}_i , setting such terms associated with small singular values to (nearly) 0 is similar to performing PCA, where we only use the \vec{v}_i associated with the largest few singular values. Thus our conclusion is — ridge regression behaves qualitatively similar to a soft form of PCA.

4.4 Tikhonov Regression

Recall that our earlier augmented system

$$\begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \vec{x} = \begin{bmatrix} \vec{y} \\ \vec{0} \end{bmatrix} \quad (4.58)$$

which had full column rank, tried to find a \vec{x} such that $A\vec{x} \approx \vec{y}$ while $\vec{x} \approx \vec{0}$ — in other words, \vec{x} that is small. Suppose that we wanted to instead enforce that \vec{x} were close to some other vector $\vec{x}_0 \in \mathbb{R}^n$. Then we would set up the system

$$\begin{bmatrix} A \\ \sqrt{\lambda}I \end{bmatrix} \vec{x} = \begin{bmatrix} \vec{y} \\ \vec{x}_0 \end{bmatrix}. \quad (4.59)$$

This would yield the least-squares type objective function

$$\|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x} - \vec{x}_0\|_2^2. \quad (4.60)$$

The final generalization of this is to put different weights on each row of $A\vec{x} - \vec{y}$ and $\vec{x} - \vec{x}_0$. If, for example, we really want to get row i of $A\vec{x}$ close to b_i , we can multiply the squared difference $(A\vec{x} - \vec{y})_i^2$ by a large weight in the loss function, and the solutions will bias towards ensuring that $(A\vec{x} - \vec{y})_i \approx 0$. Similarly, if we really are sure that the true \vec{x} has i^{th} coordinate $(\vec{x}_0)_i$, then we can attach a large weight to the difference $(\vec{x} - \vec{x}_0)_i^2$ as well. Mathematically, this gives us the following objective function:

$$\|W_1(A\vec{x} - \vec{y})\|_2^2 + \|W_2(\vec{x} - \vec{x}_0)\|_2^2, \quad (4.61)$$

where $W_1 \in \mathbb{R}^{m \times m}$ and $W_2 \in \mathbb{R}^{n \times n}$ are diagonal matrices representing the weights. Notice how this is a generalization of ridge regression with $W_1 = I$, $W_2 = \sqrt{\lambda}I$, and $\vec{x}_0 = \vec{0}$. This general regression is called Tikhonov regression.

Theorem 87 (Tikhonov Regression)

Let $A \in \mathbb{R}^{m \times n}$, $\vec{x}_0 \in \mathbb{R}^n$, and $\vec{y} \in \mathbb{R}^m$, and let $W_1 \in \mathbb{R}^{m \times m}$ and $W_2 \in \mathbb{R}^{n \times n}$ be diagonal. Then the unique solution to the *Tikhonov regression* problem

$$\min_{\vec{x} \in \mathbb{R}^n} \left\{ \|W_1(A\vec{x} - \vec{y})\|_2^2 + \|W_2(\vec{x} - \vec{x}_0)\|_2^2 \right\} \quad (4.62)$$

is given by

$$\vec{x}^* = (A^\top W_1^2 A + W_2^2)^{-1} (A^\top W_1^2 \vec{y} + W_2^2 \vec{x}_0). \quad (4.63)$$

Proof. Left as exercise. □

This expression looks complicated, so we do a sanity-check; if $W_1 = I$, $W_2 = \sqrt{\lambda}I$, and $\vec{x}_0 = \vec{0}$, then we get exactly the ridge regression solution.

4.5 Maximum Likelihood Estimation (MLE)

Previously, we talked about incorporating side information (like $\vec{x} = \vec{x}_0$) deterministically. Now we discuss a way to incorporate probabilistic information into our model.

Namely, suppose that the rows of our A matrix are vectors $\vec{a}_1, \dots, \vec{a}_m \in \mathbb{R}^n$, and that the entries of our \vec{y} vector are $y_1, \dots, y_m \in \mathbb{R}$. Now suppose we have the probabilistic model

$$y_i = \vec{a}_i^\top \vec{x} + w_i, \quad \forall i \in \{1, \dots, m\} \quad (4.64)$$

where w_1, \dots, w_m are independent Gaussian random variables; in particular, $w_i \sim \mathcal{N}(0, \sigma_i^2)$. Or in short, we have

$$\vec{y} = A\vec{x} + \vec{w} \quad (4.65)$$

where $\vec{w} = [w_1 \ \dots \ w_m]^\top \in \mathbb{R}^m$. In this case, we say that $\vec{w} \sim \mathcal{N}(\vec{0}, \Sigma_{\vec{w}})$ where $\Sigma_{\vec{w}} \doteq \text{diag}(\sigma_1^2, \dots, \sigma_m^2)$

In this setup, the *maximum likelihood estimate* (MLE) for \vec{x} turns out to be exactly a solution to a Tikhonov regression problem. The maximum likelihood estimate is the parameter choice which makes the data most likely, in that it has the highest probability or probability density out of all choices of the parameter. It is a meaningful and popular statistical estimator; thus the fact that we can reduce its computation to a ridge regression-type problem is both interesting and useful.

Henceforth, we use p to denote probability densities, and use $p_{\vec{x}}$ to denote probability densities for a fixed value of \vec{x} . In the above model, \vec{x} is not a random variable, so it doesn't quite make formal sense to condition on it (though — spoilers! — we will soon put a probabilistic prior on it, and then it makes sense to condition).

Proposition 88 (MLE as Tikhonov Regression)

In the above probabilistic model, we have

$$\operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} p_{\vec{x}}(\vec{y}) = \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \left\| \Sigma_{\vec{w}}^{-1/2} (A\vec{x} - \vec{y}) \right\|_2^2. \quad (4.66)$$

Proof. Since the logarithm is monotonically increasing, $\operatorname{argmax}_{\vec{x}} f(\vec{x}) = \operatorname{argmax}_{\vec{x}} \log(f(\vec{x}))$ for all functions f , and so

$$\operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} p_{\vec{x}}(\vec{y}) = \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \log(p_{\vec{x}}(\vec{y})) \quad (4.67)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \log\left(\prod_{i=1}^m p_{\vec{x}}(y_i)\right) \quad (4.68)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \log(p_{\vec{x}}(y_i)) \quad (4.69)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \log\left(\frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2}\right)\right) \quad (4.70)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \left\{ \underbrace{\log \left(\frac{1}{\sqrt{2\pi\sigma_i^2}} \right)}_{\text{independent of } \vec{x}} + \log \left(\exp \left(-\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2} \right) \right) \right\} \quad (4.71)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \left\{ \log \left(\exp \left(-\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2} \right) \right) \right\} \quad (4.72)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \left\{ -\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2} \right\} \quad (4.73)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \left\{ -\frac{1}{2} \sum_{i=1}^m \frac{(y_i - \vec{a}_i^\top \vec{x})^2}{\sigma_i^2} \right\} \quad (4.74)$$

$$= \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \frac{(y_i - \vec{a}_i^\top \vec{x})^2}{\sigma_i^2} \quad (4.75)$$

$$= \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \left\| \Sigma_{\vec{w}}^{-1/2} (A\vec{x} - \vec{y}) \right\|_2^2. \quad (4.76)$$

□

4.6 Maximum A Posteriori Estimation (MAP)

Now we consider the same probabilistic model as above, except this time suppose that we believe \vec{x} is also random, in the sense that

$$x_j = \mu_j + v_j, \quad \forall j \in \{1, \dots, n\} \quad (4.77)$$

where v_1, \dots, v_n are independent Gaussian random variables; in particular $v_j \sim \mathcal{N}(0, \tau_j^2)$. Or in short, we have

$$\vec{x} = \vec{x}_0 + \vec{v} \quad (4.78)$$

where $\vec{v} = [v_1 \ \dots \ v_n]^\top \in \mathbb{R}^n$ is distributed as $\vec{v} \sim \mathcal{N}(\vec{0}, \Sigma_{\vec{v}})$, where $\Sigma_{\vec{v}} \doteq \operatorname{diag}(\tau_1^2, \dots, \tau_n^2)$.

In this setup, the maximum likelihood estimate may still be useful, but another quantity that is perhaps more relevant is the *maximum a posteriori estimate* (MAP). The MAP estimate is the value of \vec{x} which is most likely, i.e., having the highest conditional probability or conditional probability density, conditioned on the observed data. It is also a meaningful and popular statistical estimator. It turns out that we can derive a similar result as in the MLE case.

Theorem 89 (MAP as Tikhonov Regression)

In the above probabilistic model, we have

$$\operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} p(\vec{x} \mid \vec{y}) = \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \left\{ \left\| \Sigma_{\vec{w}}^{-1/2} (A\vec{x} - \vec{y}) \right\|_2^2 + \left\| \Sigma_{\vec{v}}^{-1/2} (\vec{x} - \vec{x}_0) \right\|_2^2 \right\}. \quad (4.79)$$

Proof. Using Bayes' rule and the computations from before, we have

$$\operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} p(\vec{x} \mid \vec{y}) \quad (4.80)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \log(p(\vec{x} \mid \vec{y})) \quad (4.81)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \log \left(\frac{p(\vec{y} \mid \vec{x})p(\vec{x})}{p(\vec{y})} \right) \quad (4.82)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \left\{ \log(p(\vec{y} | \vec{x})) + \log(p(\vec{x})) - \underbrace{\log(p(\vec{y}))}_{\text{independent of } \vec{x}} \right\} \quad (4.83)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \{ \log(p(\vec{y} | \vec{x})) + \log(p(\vec{x})) \} \quad (4.84)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \log \left(\left\{ \prod_{i=1}^m p(y_i | \vec{x}) \right\} \cdot \left\{ \prod_{j=1}^n p(x_j) \right\} \right) \quad (4.85)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \left\{ \sum_{i=1}^m \log(p(y_i | \vec{x})) + \sum_{j=1}^n \log(p(x_j)) \right\} \quad (4.86)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \left\{ \sum_{i=1}^m \log \left(\frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left(-\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2} \right) \right) + \sum_{j=1}^n \log \left(\frac{1}{\sqrt{2\pi\tau_j^2}} \exp \left(-\frac{(x_j - (\vec{x}_0)_j)^2}{2\tau_j^2} \right) \right) \right\} \quad (4.87)$$

$$= \operatorname{argmax}_{\vec{x} \in \mathbb{R}^n} \left\{ \sum_{i=1}^m \left(-\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{2\sigma_i^2} \right) + \sum_{j=1}^n \left(-\frac{(x_j - (\vec{x}_0)_j)^2}{2\tau_j^2} \right) \right\} \quad (4.88)$$

$$= \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \left\{ \sum_{i=1}^m \left(\frac{(y_i - \vec{a}_i^\top \vec{x})^2}{\sigma_i^2} \right) + \sum_{j=1}^n \left(\frac{(x_j - (\vec{x}_0)_j)^2}{\tau_j^2} \right) \right\} \quad (4.89)$$

$$= \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \left\{ \left\| \Sigma_{\vec{w}}^{-1/2} (A\vec{x} - \vec{y}) \right\|_2^2 + \left\| \Sigma_{\vec{v}}^{-1/2} (\vec{x} - \vec{x}_0) \right\|_2^2 \right\} \quad (4.90)$$

as desired. □

Chapter 5

Convexity

Relevant sections of the textbooks:

- [1] Chapter 4.
- [2] Chapter 8.

5.1 Convex Sets

5.1.1 Basics

First, we want to define a special type of linear combination called a *convex combination*.

Definition 90 (Convex Combination)

Let $\vec{x}_1, \dots, \vec{x}_k \in \mathbb{R}^n$. The sum

$$\vec{x} = \sum_{i=1}^k \theta_i \vec{x}_i \tag{5.1}$$

is a *convex combination* of $\vec{x}_1, \dots, \vec{x}_k$ if each $\theta_i \geq 0$ and $\sum_{i=1}^k \theta_i = 1$.

We can think of each θ_i as a weight on the corresponding \vec{x}_i . Since they are non-negative numbers which sum to 1, we can also interpret them as probabilities.

Definition 91 (Convex Set)

Let $C \subseteq \mathbb{R}^n$. We say that C is a *convex set* if it is closed under convex combinations: for all $\vec{x}_1, \vec{x}_2 \in C$ and all $\theta \in [0, 1]$, we have $\theta \vec{x}_1 + (1 - \theta) \vec{x}_2 \in C$.

Geometrically, a set C is convex if for every two points $\vec{x}_1, \vec{x}_2 \in C$, the line segment $\{\theta \vec{x}_1 + (1 - \theta) \vec{x}_2 \mid \theta \in [0, 1]\}$ is contained in C . This means that, for example, the midpoint between \vec{x}_1 and \vec{x}_2 , i.e., $\frac{1}{2} \vec{x}_1 + \frac{1}{2} \vec{x}_2$, is contained in C , as well as the point $\frac{1}{3}$ of the way from \vec{x}_1 to \vec{x}_2 , i.e., $\frac{2}{3} \vec{x}_1 + \frac{1}{3} \vec{x}_2$, etc. More generally, as we vary θ , we go along the line segment connecting \vec{x}_1 and \vec{x}_2 .




Figure 5.1: Two sets $C_1, C_2 \subseteq \mathbb{R}^2$. C_1 is not convex, but C_2 is. To visualize the behavior of the line segments, we also plot a point on each line segment along with its associated θ .

Algebraically, a set C is convex if for any $\vec{x}_1, \dots, \vec{x}_k \in C$, any convex combination of $\vec{x}_1, \dots, \vec{x}_k$ is contained in C .

One way to generate a convex set from any (possibly non-convex) set, including finite and infinite sets, is to take its convex hull.

Definition 92 (Convex Hull)

Let $S \subseteq \mathbb{R}^n$ be a set. The *convex hull* of S , denoted $\text{conv}(S)$, is the set of all convex combinations of points in S , i.e.,

$$\text{conv}(S) = \left\{ \sum_{i=1}^k \theta_i \vec{x}_i \mid k \in \mathbb{N}, \theta_1, \dots, \theta_k \geq 0, \sum_{i=1}^k \theta_i = 1, \vec{x}_1, \dots, \vec{x}_k \in S \right\}. \quad (5.2)$$

Here are some properties of the convex hull; the proof is left as an exercise.

Proposition 93

Let $S \subseteq \mathbb{R}^n$ be a set.

- (a) $\text{conv}(S)$ is a convex set.
- (b) $\text{conv}(S)$ is the minimal convex set which contains S , i.e.,

$$\text{conv}(S) = \bigcap_{\substack{C \supseteq S \\ C \text{ is a convex set}}} C. \quad (5.3)$$

Thus if S is convex then $\text{conv}(S) = S$.

- (c) $\text{conv}(S)$ is the union of convex hulls of all finite subsets of S , i.e.,

$$\text{conv}(S) = \bigcup_{\substack{A \subseteq S \\ A \text{ is a finite set}}} \text{conv}(A). \quad (5.4)$$

Actually, the last statement can be strengthened to a separate, more quantitative result, which gives a fundamental characterization of convex sets.

Theorem 94 (Carathéodory's Theorem)

Let $S \subseteq \mathbb{R}^n$ be a set. Then $\text{conv}(S)$ is the union of convex hulls of all finite subsets of S of size at most $n + 1$, i.e.,

$$\text{conv}(S) = \bigcup_{\substack{A \subseteq S \\ |A| \leq n+1}} \text{conv}(A). \quad (5.5)$$

The proof of this theorem is left as an exercise; interested students can reference the proof in Bertsekas [5, Proposition B.6], for example.

Below, we visualize the convex hull of a *finite set* S . By the above proposition, the convex hull of an infinite set S' is the union of convex hulls of all finite subsets of S' .

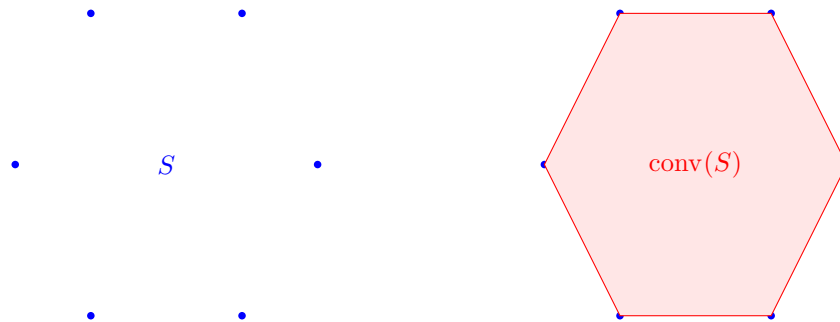


Figure 5.2: A finite set S and its convex hull.

The following content is optional/out of scope for this semester. Regardless, it may be helpful to read it to gain context, or get a deeper understanding of various results.

5.1.2 (OPTIONAL) Conic Hull, Affine Hull, and Relative Interior

As stated above, given a set S that is not necessarily convex, taking arbitrary convex combinations of vectors in S generates the convex hull $\text{conv}(S)$ of S , which is the smallest convex set containing S . Below, we present two other methods of generating convex sets containing S that look similar to the definition of the convex hull. Each method has its own geometric interpretation.

Definition 95 (Conic Hull)

Let $S \subseteq \mathbb{R}^n$ be a set. The *conic hull* of S , denoted $\text{conic}(S)$, is defined as the set of *conic combinations* of vectors in S , i.e., linear combinations of vectors in S with non-negative coefficients:

$$\text{conic}(S) = \left\{ \sum_{i=1}^k \theta_i \vec{x}_i \mid k \in \mathbb{N}, \theta_1, \dots, \theta_k \geq 0, \vec{x}_1, \dots, \vec{x}_k \in S \right\}. \quad (5.6)$$

Geometrically, the conic hull of a set S is the set of all rays from the origin that pass through $\text{conv}(S)$.

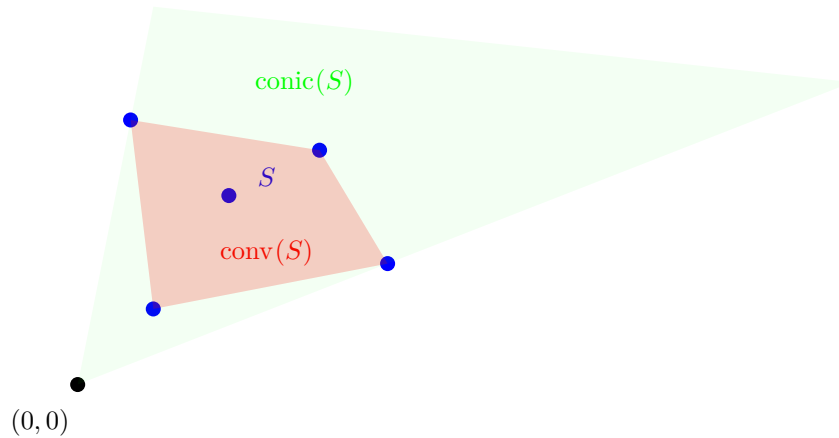


Figure 5.3: A finite set S and its convex and conic hulls.

Definition 96 (Affine Set)

Let $S \subseteq \mathbb{R}^n$ be a set. We say that S is an *affine set* if it is closed under affine combinations: for each $\vec{x}_1, \vec{x}_2 \in S$, and any $\theta \in \mathbb{R}$, we have $\theta\vec{x}_1 + (1 - \theta)\vec{x}_2 \in S$.

Note the difference between affine and convex sets. In the latter, θ is restricted to $[0, 1]$. Geometrically this restriction corresponds to the (finite) line segment connecting \vec{x}_1 and \vec{x}_2 being contained in S . In the former, however, θ can be any real number, corresponding to the whole (infinite) line connecting \vec{x}_1 and \vec{x}_2 being contained in S .

Note that an affine set is a translation of a subspace. This intuition is one of the most helpful ways to understand affine sets.

Proposition 97

For a set $A \subseteq \mathbb{R}^n$, define the translation $A + \vec{x} \doteq \{\vec{a} + \vec{x} \mid \vec{a} \in A\}$.

- (a) Let $S \subseteq \mathbb{R}^n$ be a nonempty affine set. Then there is a subspace $U \subseteq \mathbb{R}^n$ such that, for any $\vec{x} \in S$, we have $S = U + \vec{x}$.
- (b) For any subspace $U \subseteq \mathbb{R}^n$ and vector $\vec{x} \in \mathbb{R}^n$, the set $U + \vec{x}$ is an affine set.

Proof.

- (a) Let $\vec{x} \in S$ be any vector in S , and define $U := S + (-\vec{x}) = \{\vec{s} - \vec{x} \mid \vec{s} \in S\}$. We claim that U is a subspace. Indeed, since $\vec{x} \in S$, we see that $\vec{0} \in U = S + (-\vec{x})$. We show that U is closed under addition. Let $\vec{u}_1, \vec{u}_2 \in U$. By definition of U , there exist $\vec{s}_1, \vec{s}_2 \in S$ such that $\vec{u}_1 = \vec{s}_1 - \vec{x}$ and $\vec{u}_2 = \vec{s}_2 - \vec{x}$. Then

$$\vec{u}_1 + \vec{u}_2 = \vec{s}_1 - \vec{x} + \vec{s}_2 - \vec{x} \tag{5.7}$$

$$= \left[2 \left(\frac{\vec{s}_1 + \vec{s}_2}{2} \right) + (1 - 2)\vec{x} \right] - \vec{x}. \tag{5.8}$$

Now because S is affine it is convex, so $\frac{\vec{s}_1 + \vec{s}_2}{2} \in S$. As an affine combination of elements in S , we have $2 \left(\frac{\vec{s}_1 + \vec{s}_2}{2} \right) + (1 - 2)\vec{x} \in S$. Thus we have $\vec{u}_1 + \vec{u}_2 \in S + (-\vec{x}) = U$, so that U is closed under vector addition.

To show that U is closed under scalar multiplication, let $\alpha \in \mathbb{R}$ and $\vec{u} \in U$. By definition of U , there exists

$\vec{s} \in S$ such that $\vec{u} = \vec{s} - \vec{x}$. Then

$$\alpha \vec{u} = \alpha(\vec{s} - \vec{x}) \quad (5.9)$$

$$= [\alpha \vec{s} + (1 - \alpha)\vec{x}] - \vec{x}. \quad (5.10)$$

Since S is affine it is convex, so $\alpha \vec{s} + (1 - \alpha)\vec{x} \in S$. Thus $\alpha \vec{u} \in S + (-\vec{x}) = U$, so U is closed under scalar multiplication. We have shown that U is closed under linear combinations and contains $\vec{0}$, so U is a subspace and the claim is proved.

(b) Let $\alpha \in \mathbb{R}$ and let $\vec{s}_1, \vec{s}_2 \in U + \vec{x}$. By definition of U , there exist $\vec{u}_1, \vec{u}_2 \in S$ such that $\vec{s}_1 = \vec{u}_1 + \vec{x}$ and $\vec{s}_2 = \vec{u}_2 + \vec{x}$. Then

$$\alpha \vec{s}_1 + (1 - \alpha)\vec{s}_2 = \alpha(\vec{u}_1 + \vec{x}) + (1 - \alpha)(\vec{u}_2 + \vec{x}) \quad (5.11)$$

$$= [\alpha \vec{u}_1 + (1 - \alpha)\vec{u}_2] + \vec{x}. \quad (5.12)$$

Since U is a subspace, $\alpha \vec{u}_1 + (1 - \alpha)\vec{u}_2 \in U$. Thus, from above, we have $\alpha \vec{s}_1 + (1 - \alpha)\vec{s}_2 = [\alpha \vec{u}_1 + (1 - \alpha)\vec{u}_2] + \vec{x} \in U + \vec{x}$. We have shown that $U + \vec{x}$ is closed under affine combinations, so it is an affine set. □

Definition 98 (Affine Hull)

Let $S \subseteq \mathbb{R}^n$ be a set. The *affine hull* of S , denoted $\text{aff}(S)$, is defined as the set of *affine combinations* of vectors in S , i.e., linear combinations of vectors in S with coefficients which sum to 1:

$$\text{aff}(S) = \left\{ \sum_{i=1}^k \theta_i \vec{x}_i \mid k \in \mathbb{N}, \theta_1, \dots, \theta_k \in \mathbb{R}, \sum_{i=1}^k \theta_i = 1, \vec{x}_1, \dots, \vec{x}_k \in S \right\}. \quad (5.13)$$

Here are some properties of the affine hull; the proof is left as an exercise.

Proposition 99

Let $S \subseteq \mathbb{R}^n$ be a set.

- (a) $\text{aff}(S)$ is an affine set.
- (b) $\text{aff}(S)$ is the minimal affine set which contains S , i.e.,

$$\text{aff}(S) = \bigcap_{\substack{C \supseteq S \\ C \text{ is an affine set}}} C. \quad (5.14)$$

Thus if S is affine then $\text{aff}(S) = S$.

- (c) $\text{aff}(S)$ is the union of affine hulls of all finite subsets of S , i.e.,

$$\text{aff}(S) = \bigcup_{\substack{A \subseteq S \\ A \text{ is a finite set}}} \text{aff}(A). \quad (5.15)$$

We can actually get an elementary refinement of (c) above.

Corollary 100. *Let $S \subseteq \mathbb{R}^n$ be a set, and let $\text{aff}(S)$ be the translation of a linear subspace of dimension $d \leq n$. Then $\text{aff}(S)$ is the union of affine hulls of all finite subsets of S of size at most d , i.e.,*

$$\text{aff}(S) = \bigcup_{\substack{A \subseteq S \\ |A| \leq d}} \text{aff}(A). \quad (5.16)$$

Proof. Suppose that $\text{aff}(S) = U + \vec{x}$ where $U \subseteq \mathbb{R}^n$ is a subspace of dimension d . We prove both subset relations, i.e., $A \subseteq B$ and $B \subseteq A$ implies $A = B$.

We show the quicker subset inequality first. We have from earlier results that

$$\text{aff}(S) = \bigcup_{\substack{A \subseteq S \\ A \text{ is a finite set}}} \text{aff}(A) \supseteq \bigcup_{\substack{A \subseteq S \\ |A| \leq d}} \text{aff}(A). \quad (5.17)$$

Towards showing the reverse inequality, let $\vec{s}_1, \dots, \vec{s}_d$ be elements of S such that, if we define $\vec{u}_i \doteq \vec{s}_i - \vec{x}$, then $\vec{u}_1, \dots, \vec{u}_d$ is a basis for U . (Such \vec{s}_i have to exist; if they don't, then there are no $\vec{s}_1, \dots, \vec{s}_d$ whose translates by \vec{x} span U , a contradiction with the definition of $\text{aff}(S) = U + \vec{x}$). Now taking $A = \{\vec{s}_1, \dots, \vec{s}_d\}$, we see that $\text{aff}(A) = \text{aff}(S)$. Thus we have

$$\text{aff}(S) = \text{aff}(\{\vec{s}_1, \dots, \vec{s}_d\}) \subseteq \bigcup_{\substack{A \subseteq S \\ |A| \leq d}} \text{aff}(A). \quad (5.18)$$

Therefore we have

$$\text{aff}(S) = \bigcup_{\substack{A \subseteq S \\ |A| \leq d}} \text{aff}(A), \quad (5.19)$$

as desired. □

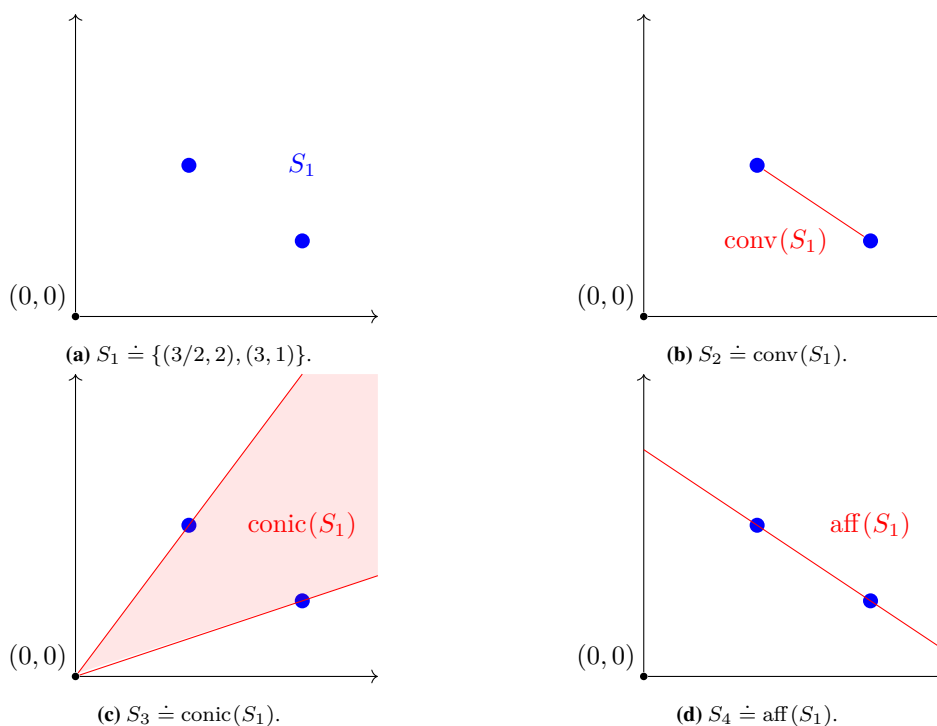


Figure 5.4: (a) The set $S_1 \doteq \{(3/2, 2), (3, 1)\}$ of two points in \mathbb{R}^2 ; (b) The convex hull $S_2 \doteq \text{conv}(S_1)$ of S_1 , which is the closed line segment connecting the two points in S_1 ; (c) The conic hull $S_3 \doteq \text{conic}(S_1)$ of S_1 , which is the union of all rays passing through S_1 ; (d) The affine hull $S_4 \doteq \text{aff}(S_1)$ of S_1 , which is the infinite line connecting the two points in S_1 . Note that we also have $S_3 = \text{conic}(S_2)$ and $S_4 = \text{aff}(S_2)$; this relationship can be shown to hold in general from definitions.

Next, given a set $S \subseteq \mathbb{R}^n$, we sometimes wish to distinguish points that lie on the “boundary” of S from points that lie in the “interior” of S . For instance, consider the set $S \doteq [0, 1] \subseteq \mathbb{R}$. In this case, 0 and 1 lie on the “boundary” of S , since they are infinitely close to both points inside S and points outside S . On the other hand, $\frac{1}{10}$ lies in the interior of S , since all points within a sufficiently small distance lie in S . Although 0 and 1 can both be geometrically interpreted as points on the boundary of $S = [0, 1]$, note that $0 \in S$ while $1 \notin S$. In general, a set may contain either all, some, or none of the points on its boundary.

Below, we formalize the notion of interior points¹.

Definition 101 (Interior)

Let $S \subseteq \mathbb{R}^n$ and let $\vec{x} \in \mathbb{R}^n$.

- (a) (*Open ball.*) Let $r > 0$. We call $N_r(\vec{x}) \doteq \{\vec{y} \in \mathbb{R}^n \mid \|\vec{y} - \vec{x}\|_2 < r\}$ the *open ball* in \mathbb{R}^n of radius r centered at \vec{x} .
- (b) (*Interior.*) We say that \vec{x} is an *interior point* of S when there exists some $r > 0$ such that $N_r(\vec{x}) \subseteq S$.^a The set of all interior points of S is called the *interior* of S and denoted $\text{int}(S)$.

^aIn the definition for the interior point, it does not matter whether we use \subseteq or \subset . Think about why; this is a good exercise to internalize the definitions.

¹The definitions provided below can be generalized to spaces more abstract than \mathbb{R}^n or even general finite-dimensional vector spaces, such as metric or topological spaces.

In words, given a set $S \subseteq \mathbb{R}^n$, we say that $\vec{x} \in \mathbb{R}^n$ is an interior point of S if it is contained inside an open ball in \mathbb{R}^n that is in turn entirely contained in S . A mental picture is provided in Figure 5.5.




Figure 5.5: The vector \vec{x} is an interior point of the set $S \subseteq \mathbb{R}^2$, since there exists a 2-dimensional ball (red) centered at \vec{x} that is contained in S .

Some sets may represent geometric shapes embedded in Euclidean space of strictly higher dimension, and therefore must have empty interior. As an example, consider the set S_2 defined in Figure 5.4(b), which connects the points $(3/2, 2)$ and $(3, 1)$ in \mathbb{R}^2 . This is a one-dimensional line segment embedded in an Euclidean space of dimension 2. Indeed, if one claims that a point in S_2 , say, the midpoint $(9/4, 3/2)$, is an interior point of S_2 , then one would have to show the existence of a two-dimensional open ball centered at $(9/4, 3/2)$ that lies entirely in S_2 . This is impossible, since S_1 is a one-dimensional line segment, and so S_2 has empty interior.

However, it may still be geometrically meaningful to classify points in such a set as points on the “edge” of the set, or points “inside” the set. In the context of the line segment S_2 , the end points $S_1 = \{(3/2, 2), (3, 1)\}$ appear at the “edge” of S_2 , while the remaining points $S_2 \setminus S_1$ are located “inside” S_2 . As we explained above, this cannot be captured using the definitions of interior points and the interior presented in Definition 101. Roughly speaking, this is because $S_2 \setminus S_1$ can only be considered points “inside” the line segment S_2 from a one-dimensional perspective, e.g., relative to the line $S_4 = \text{aff}(S_1) = \text{aff}(S_2)$ that contains S_2 . This motivates the following definition of *relative interior*, provided below.

Definition 102 (Relative Interior Points, Relative Interior)

Let $S \subseteq \mathbb{R}^n$ be a set, and let $\vec{x} \in S$. We say that \vec{x} is a *relative interior point* of S when there exists some $r > 0$ such that $N_r(\vec{x}) \cap \text{aff}(S) \subseteq S$. The relative interior of S is the set of all relative interior points of S , and is denoted $\text{relint}(S)$.

In words, given a set $S \subseteq \mathbb{R}^n$, we say that $\vec{x} \in \mathbb{R}^n$ is a relative interior point of S if it is contained inside an open ball in \mathbb{R}^n whose intersection with $\text{aff}(S)$ is entirely contained in S .




Figure 5.6: A similar setup to Figure 5.4. Here $S_1 \doteq \{(3/2, 2), (3, 1)\}$, $S_2 \doteq \text{conv}(S_1) = \{\theta(3/2, 2) + (1-\theta)(3, 1) \mid \theta \in [0, 1]\}$, and $S_3 \doteq \text{aff}(S_1) = \text{aff}(S_2) = \{\theta(3/2, 2) + (1-\theta)(3, 1) \mid \theta \in \mathbb{R}\}$. Thus $\text{relint}(S_2) = S_2 \setminus S_1$. In other words, S_2 is the line segment connecting $(3/2, 2)$ and $(3, 1)$, $S_3 = \text{aff}(S_2)$ is the extension of S_2 into a line, and $\text{relint}(S_2)$ is the open (i.e., excluding the endpoints) line segment connecting $(3/2, 2)$ and $(3, 1)$. This illustrates the description of the relative interior of a set as its interior when viewed as a subset of its own affine hull.




Figure 5.7: A set S and its affine hull. While the interior of S is the empty set, its relative interior is nonempty.

Next, we use the concept of relative interior to characterize *strictly convex sets*.

Definition 103 (Strictly Convex Sets)

Let $C \subseteq \mathbb{R}^n$ be a set. We say that C is a *strictly convex set* if for every $\vec{x}_1, \vec{x}_2 \in C$ and each $\theta \in (0, 1)$, we have $\theta\vec{x}_1 + (1-\theta)\vec{x}_2 \in \text{relint}(C)$.




Figure 5.8: *Left:* a strictly convex set S_1 . *Middle:* a convex set S_2 which is not strictly convex. *Right:* A non-convex set S_3 . All three sets are defined to include their boundaries. In particular, S_2 is not strictly convex because some sections of its boundary consist of line segments. For any two points along the same line segment, each convex combination of these points will lie on the boundary of S_2 .

The above content is optional/out of scope for this semester, but now we resume the required/in scope content.

5.1.3 Hyperplane and Half-Spaces

Definition 104 (Hyperplane)

Let $\vec{a}, \vec{x}_0 \in \mathbb{R}^n$ and $b \in \mathbb{R}$. A *hyperplane* is a set of the form

$$\{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top \vec{x} = b\} \quad (5.20)$$

or, equivalently, a set of the form

$$\{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) = 0\}. \quad (5.21)$$

The equations $\vec{a}^\top \vec{x} = b$ and $\vec{a}^\top (\vec{x} - \vec{x}_0) = 0$ are connected, because if we define $b = \vec{a}^\top \vec{x}_0$, then the second equation resolves to the first equation; and if we take \vec{x}_0 to be any vector such that $\vec{a}^\top \vec{x}_0 = b$, then the first equation resolves to the second equation.

Example 105. Hyperplanes are convex. Consider a hyperplane $H \doteq \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top \vec{x} = b\}$. Let $\vec{x}_1, \vec{x}_2 \in H$ and $\theta \in [0, 1]$. Then

$$\vec{a}^\top (\theta \vec{x}_1 + (1 - \theta) \vec{x}_2) = \theta \vec{a}^\top \vec{x}_1 + (1 - \theta) \vec{a}^\top \vec{x}_2 \quad (5.22)$$

$$= \theta b + (1 - \theta) b \quad (5.23)$$

$$= b \quad (5.24)$$

so $\theta \vec{x}_1 + (1 - \theta) \vec{x}_2 \in H$. Thus H is convex.

To show that a set C is convex, we need to show that for *every* $\vec{x}_1, \vec{x}_2 \in C$ and *every* $\theta \in [0, 1]$, that $\theta \vec{x}_1 + (1 - \theta) \vec{x}_2 \in C$.

To show that C is *not* convex, we just need to come up with *one* choice of $\vec{x}_1, \vec{x}_2 \in C$ and *one* $\theta \in [0, 1]$ such that $\theta\vec{x}_1 + (1 - \theta)\vec{x}_2 \notin C$. Note that even if C is non-convex, there could be *some* choices of $\vec{x}_1, \vec{x}_2 \in C, \theta \in [0, 1]$ such that $\theta\vec{x}_1 + (1 - \theta)\vec{x}_2 \in C$; but if C is non-convex, there is at least one choice of $\vec{x}_1, \vec{x}_2 \in C, \theta \in [0, 1]$ such that $\theta\vec{x}_1 + (1 - \theta)\vec{x}_2 \notin C$.

Definition 106 (Half-Space)

Let $\vec{a}, \vec{x}_0 \in \mathbb{R}^n$ and $b \in \mathbb{R}$. A *positive half-space* is a set of the form

$$\{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top \vec{x} \geq b\} \quad \text{or} \quad \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) \geq 0\}. \quad (5.25)$$

A *negative half-space* is a set of the form

$$\{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top \vec{x} \leq b\} \quad \text{or} \quad \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) \leq 0\}. \quad (5.26)$$


The mental picture we have for these hyperplanes and half-spaces is the following. Let $\vec{x}_0 \in \mathbb{R}^n$ and define

$$H \doteq \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) = 0\} \quad (5.27)$$


$$H_+ \doteq \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) \geq 0\} \quad (5.28)$$

$$H_- \doteq \{\vec{x} \in \mathbb{R}^n \mid \vec{a}^\top (\vec{x} - \vec{x}_0) \leq 0\}. \quad (5.29)$$


Then the alignment of these objects looks like the following:



In words, the positive and negative half-spaces partition \mathbb{R}^n . Looking at some individual vectors, say $\vec{x}_1 \in H_-$ and $\vec{x}_2 \in H_+$, we have the picture



If we draw lines connecting \vec{x}_0 with \vec{x}_1 and \vec{x}_2 , they are not themselves representations of \vec{x}_1 and \vec{x}_2 , unless $\vec{x}_0 = \vec{0}$. Instead, they are representations of the *displacements* of \vec{x}_1 and \vec{x}_2 from \vec{x}_0 . Thus, we see the following picture:



And this gives us a clearer understanding of what's going on — $\vec{x}_1 - \vec{x}_0$ forms an obtuse angle with \vec{a} , indicating a negative dot product, whereas $\vec{x}_2 - \vec{x}_0$ forms an acute angle with \vec{a} , indicating a positive dot product. And this is how H_+ and H_- are computed.

This allows us to consider what it means for a hyperplane to separate two sets. It means that for every vector in the first set, the dot product is non-positive, and for every vector in the second set, the dot product is non-negative.

Example 107 (Set of PSD Matrices is Convex). Consider \mathbb{S}_+^n , the set of all symmetric positive semidefinite (PSD) matrices. We want to show that \mathbb{S}_+^n is convex. Take $A_1, A_2 \in \mathbb{S}_+^n$ and $\theta \in [0, 1]$. We want to show that $\theta A_1 + (1-\theta)A_2 \in \mathbb{S}_+^n$.

One of the ways to tell if a matrix A is PSD is to check whether $\vec{x}^\top A \vec{x} \geq 0$ for all $\vec{x} \in \mathbb{R}^n$. Checking this for our convex combination, we get

$$\vec{x}^\top (\theta A_1 + (1-\theta)A_2) \vec{x} = \theta \underbrace{\vec{x}^\top A_1 \vec{x}}_{\geq 0} + (1-\theta) \underbrace{\vec{x}^\top A_2 \vec{x}}_{\geq 0} \quad (5.30)$$

$$\geq 0. \quad (5.31)$$

Note that it is possible to come up with linear combinations of PSD matrices that are not PSD; indeed, $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and

$\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$ are PSD, yet their difference $\begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$ is not PSD. But all convex combinations of PSD matrices are PSD, as we have confirmed above.

Theorem 108 (Separating Hyperplane Theorem)

Let $C, D \subseteq \mathbb{R}^n$ be two nonempty *disjoint* convex sets, i.e., $C \cap D = \emptyset$. Then there exists a hyperplane that separates C and D , i.e., there exists $\vec{a}, \vec{x}_0 \in \mathbb{R}^n$ such that^a

$$\vec{a}^\top (\vec{x} - \vec{x}_0) \geq 0, \quad \forall \vec{x} \in C \quad (5.32)$$

$$\vec{a}^\top (\vec{x} - \vec{x}_0) \leq 0, \quad \forall \vec{x} \in D. \quad (5.33)$$


Moreover, if C is closed (containing its boundary points) and D is closed and bounded, then there exists a hyperplane that separates C and D without intersecting either set, i.e., there exists $\vec{a}, \vec{x}_0 \in \mathbb{R}^n$ such that

$$\vec{a}^\top (\vec{x} - \vec{x}_0) > 0, \quad \forall \vec{x} \in C \quad (5.34)$$

$$\vec{a}^\top (\vec{x} - \vec{x}_0) < 0, \quad \forall \vec{x} \in D. \quad (5.35)$$

^aBy defining $b \doteq \vec{a}^\top \vec{x}_0$, one can express $\vec{a}^\top (\vec{x} - \vec{x}_0)$ as $\vec{a}^\top \vec{x} - b$ if desired.

The mental picture we want to have is the following.




Proof. We prove the part of the theorem statement in the case where C is closed and bounded and D is closed.

Even though our theorem statement concerns existence of such \vec{a} and \vec{x}_0 , we will prove it by construction, i.e., we will construct a \vec{a} and \vec{x}_0 which separate C and D . This proof strategy is very powerful and will show up frequently.


Since C and D are disjoint, any points in C and D are separated by some positive distance; since they are compact, this distance has a finite lower bound.² Define

$$\text{dist}(C, D) \doteq \min_{\substack{\vec{c} \in C \\ \vec{d} \in D}} \|\vec{c} - \vec{d}\|_2. \quad (5.36)$$

Note that $\text{dist}(C, D) > 0$, and there exists some $c \in C$ and $d \in D$ such that $\|\vec{c} - \vec{d}\|_2 = \text{dist}(C, D)$.³



This signals that we want $\vec{c} - \vec{d}$ to be the *normal* vector of our hyperplane — that is, our \vec{a} vector. To find the other point \vec{x}_0 which the hyperplane passes through, we can just have it pass through the midpoint of \vec{c} and \vec{d} , i.e., $\frac{\vec{c} + \vec{d}}{2}$. This gives the following diagram.




Thus our proposed hyperplane has \vec{a} and \vec{x}_0 equal to

$$\vec{a} = \vec{c} - \vec{d}, \quad \vec{x}_0 = \frac{\vec{c} + \vec{d}}{2}. \quad (5.37)$$

It yields the following picture, where the hyperplane is a dotted line.

²Proving this requires some mathematical analysis and is out of scope of the course.

³Same as the above footnote. The fact that C is closed and bounded and D is closed will not be used from this point onwards.



Notice that there are many separating hyperplanes, such as the one discussed before the theorem. But we just need to prove that this hyperplane separates C and D .

The equation for this hyperplane is

$$\vec{a}^\top (\vec{x} - \vec{x}_0) = (\vec{c} - \vec{d})^\top \left(\vec{x} - \frac{\vec{c} + \vec{d}}{2} \right) \quad (5.38)$$

$$= (\vec{c} - \vec{d})^\top \vec{x} - \frac{(\vec{c} - \vec{d})^\top (\vec{c} + \vec{d})}{2} \quad (5.39)$$

$$= (\vec{c} - \vec{d})^\top \vec{x} - \frac{\vec{c}^\top \vec{c} - \vec{d}^\top \vec{d}}{2} \quad (5.40)$$

$$= (\vec{c} - \vec{d})^\top \vec{x} - \frac{\|\vec{c}\|_2^2 - \|\vec{d}\|_2^2}{2}. \quad (5.41)$$

Thus the given hyperplane is also available in (\vec{a}, b) form as

$$\vec{a} = \vec{c} - \vec{d}, \quad b = \frac{\|\vec{c}\|_2^2 - \|\vec{d}\|_2^2}{2}. \quad (5.42)$$

Now we prove that it actually separates \vec{c} and \vec{d} . Define $f: \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$f(\vec{x}) = (\vec{c} - \vec{d})^\top \left(\vec{x} - \frac{\vec{c} + \vec{d}}{2} \right). \quad (5.43)$$

For the sake of contradiction, suppose there exists $\vec{u} \in D$ such that $f(\vec{u}) \geq 0$. We can write

$$0 \leq f(\vec{u}) \quad (5.44)$$

$$= (\vec{c} - \vec{d})^\top \left(\vec{u} - \frac{\vec{c} + \vec{d}}{2} \right) \quad (5.45)$$

$$= (\vec{c} - \vec{d})^\top \left(\vec{u} - \vec{d} - \frac{\vec{c} - \vec{d}}{2} \right) \quad (5.46)$$


$$= (\vec{c} - \vec{d})^\top (\vec{u} - \vec{d}) - \frac{(\vec{c} - \vec{d})^\top (\vec{c} - \vec{d})}{2} \quad (5.47)$$

$$= (\vec{c} - \vec{d})^\top (\vec{u} - \vec{d}) - \frac{1}{2} \|\vec{c} - \vec{d}\|_2^2. \quad (5.48)$$

Thus


$$0 \leq (\vec{c} - \vec{d})^\top (\vec{u} - \vec{d}) - \frac{1}{2} \|\vec{c} - \vec{d}\|_2^2 < (\vec{c} - \vec{d})^\top (\vec{u} - \vec{d}). \quad (5.49)$$

This means that $\vec{c} - \vec{d}$ and $\vec{u} - \vec{d}$ form an *acute* angle. It also means that $\vec{u} \neq \vec{d}$, since otherwise the dot product would be 0. Going back to our picture, this means that \vec{u} would have to be positioned similarly to the following:



At least from the diagram, it seems hard to imagine a $\vec{u} \in D$ such that $\vec{u} - \vec{d}$ and $\vec{c} - \vec{d}$ form an acute angle. Namely, any vector $\vec{x} \in \mathbb{R}^n$ (of reasonably small norm, such as the \vec{u} in the figure) such that $\vec{x} - \vec{d}$ and $\vec{c} - \vec{d}$ form an acute angle, seems to be closer to \vec{c} than \vec{d} is to \vec{c} .

Why do we need the “reasonably small norm” condition? Consider the following possible \vec{x} :



Certainly, this \vec{x} is farther from \vec{c} than \vec{d} is, and so no contradiction would be derived.

If we can prove that our \vec{u} , which we assume is in D , is closer to \vec{c} than \vec{d} is, then we can derive a contradiction with the fact that \vec{d} is the closest vector in D to \vec{c} . But we can't prove this for our \vec{u} directly, because $\|\vec{u} - \vec{d}\|_2$ may be large as in the above figure, so instead we take another vector \vec{x} which is close to \vec{d} , where the displacement between \vec{x} and \vec{d} points in the direction of \vec{u} . We will show that this \vec{x} is in D yet is closer to \vec{c} than \vec{d} is, thus deriving a contradiction.

Here are the details. Let $\vec{p}: [0, 1] \rightarrow \mathbb{R}^n$ trace out the line from \vec{d} to \vec{u} ; namely, let $\vec{p}(t) = \vec{d} + t(\vec{u} - \vec{d}) = t\vec{u} + (1-t)\vec{d}$. Since $\vec{u}, \vec{d} \in D$ by assumption, and D is convex, we have that $\vec{p}(t) \in D$ for all $t \in [0, 1]$. Now we see that

$$\begin{aligned} \|\vec{p}(t) - \vec{c}\|_2^2 &= \|\vec{d} + t(\vec{u} - \vec{d}) - \vec{c}\|_2^2 \\ &= \|(\vec{d} - \vec{c}) + t(\vec{u} - \vec{d})\|_2^2 \\ &= \|\vec{d} - \vec{c}\|_2^2 + 2t(\vec{d} - \vec{c})^\top (\vec{u} - \vec{d}) + t^2 \|\vec{u} - \vec{d}\|_2^2 \\ &= \|\vec{c} - \vec{d}\|_2^2 - 2t(\vec{c} - \vec{d})^\top (\vec{u} - \vec{d}) + t^2 \|\vec{u} - \vec{d}\|_2^2. \end{aligned}$$

We want to show that there exists t such that $\|\bar{p}(t) - \bar{c}\|_2^2 < \|\bar{c} - \bar{d}\|_2^2$, i.e.,

$$-2t(\bar{c} - \bar{d})^\top (\bar{u} - \bar{d}) + t^2 \|\bar{u} - \bar{d}\|_2^2 < 0, \quad (5.50)$$

i.e.,

$$2 \underbrace{(\bar{c} - \bar{d})^\top (\bar{u} - \bar{d})}_{>0} - t \|\bar{u} - \bar{d}\|_2^2 > 0. \quad (5.51)$$

Now for all $0 < t < \frac{2(\bar{c} - \bar{d})^\top (\bar{u} - \bar{d})}{\|\bar{u} - \bar{d}\|_2^2}$, i.e., t small enough, the above inequality holds, so we have for this t that

$$\begin{aligned} \|\bar{p}(t) - \bar{c}\|_2^2 &= \|\bar{c} - \bar{d}\|_2^2 - 2t(\bar{c} - \bar{d})^\top (\bar{u} - \bar{d}) + t^2 \|\bar{u} - \bar{d}\|_2^2 \\ &< \|\bar{c} - \bar{d}\|_2^2. \end{aligned}$$

However, $\bar{p}(t) \in D$, a contradiction. \square

The following content is optional/out of scope for this semester. Regardless, it may be helpful to read it to gain context, or get a deeper understanding of various results.

5.1.4 (OPTIONAL) Cones

Definition 109 (Cones, Proper Cones)

Let $K \subseteq \mathbb{R}^n$.

- (a) We call K a *cone* if, for any $\vec{v} \in K$ and $\alpha \in \mathbb{R}_+^a$, we have $\alpha\vec{v} \in K$.
- (b) We call K a *convex cone* if it is both a cone and a convex set.
- (c) We call K a *pointed cone* if it contains no line through the origin, i.e., if for each nonzero $\vec{v} \in K$, there exists some $\alpha \in \mathbb{R}$ such that $\alpha\vec{v} \notin K$.
- (d) We call K a *solid cone* if it has non-empty interior, i.e., if there exists some $\vec{v} \in K$ and some $r > 0$ such that the open ball in \mathbb{R}^n of radius r centered at \vec{v} is contained in K : namely, we have $\{\vec{w} \in \mathbb{R}^n \mid \|\vec{w} - \vec{v}\|_2 < r\} \subseteq K$.
- (e) We call K a *closed cone* if it is a closed set, i.e., it contains its boundary points.
- (f) We call K a *proper cone* if it is convex, pointed, solid, and closed.

^aThat is, $\alpha \in \mathbb{R}$ and $\alpha \geq 0$

Note that non-empty cones must contain the zero vector, which corresponds to the case of taking $\alpha = 0$ in the definition of a cone.

The definition of proper cones is motivated by their connection to generalized inequalities in convex optimization, which will be discussed later in the course in the context of second-order cone programs (SOCPs) and semidefinite programming (SDPs). For this we require the above definitions to apply to a slightly broader context. We would need to replace \mathbb{R}^n with a generic vector space V and the $\|\cdot\|_2$ norm with any norm $\|\cdot\|_V$ on this vector space. In fact, for the following results to hold we additionally need to have an inner product on this vector space $\langle \cdot, \cdot \rangle_V$ that is compatible with the norm, i.e., $\langle \vec{x}, \vec{x} \rangle_V = \|\vec{x}\|_V^2$; that is, we would need an *inner product space*. One can check (as an exercise)

that \mathbb{R}^n is an example of such an inner product space, with the ℓ^2 norm $\|\cdot\|_2$ and usual inner product. Thus, in order to generalize the results introduced in this section, we would replace \mathbb{R}^n with V , replace the norm $\|\vec{x}\|_2$ with $\|\vec{x}\|_V$, and replace the inner product $\vec{x}^\top \vec{y}$ with $\langle \vec{x}, \vec{y} \rangle_V$.⁴


For now, we return to working over \mathbb{R}^n so as to not introduce additional complexity in the definitions.

Example 110. The sets

$$C_1 \doteq \{(\vec{x}, y) \in \mathbb{R}^{n+1} \mid \|\vec{x}\|_2 \leq y\} \quad (5.52)$$

$$C_2 \doteq \{(x, y) \in \mathbb{R}^2 \mid y \geq 0\} \quad (5.53)$$

are cones. If $n = 1$ then C_1 looks like:



Now, we discuss some important classes of cones.

Definition 111

(a) A set of the form

$$C \doteq \{(\vec{x}, t) \in \mathbb{R}^{n+1} \mid A\vec{x} \leq t\vec{y}, t \geq 0\} \quad (5.54)$$

is called a *polyhedral cone*, and in particular corresponds to the polyhedron $\{\vec{x} \in \mathbb{R}^n \mid A\vec{x} \leq \vec{y}\}$.

(b) A set of the form

$$C \doteq \{(\vec{x}, t) \in \mathbb{R}^{n+1} \mid \|A\vec{x} - t\vec{y}\|_2 \leq tz, t \geq 0\} \quad (5.55)$$

is called a *ellipsoidal cone*, and in particular corresponds to the ellipse $\{\vec{x} \in \mathbb{R}^n \mid \|A\vec{x} - \vec{y}\|_2 \leq z\}$.^a

^aThe ellipsoidal cone corresponding to the unit circle — which is, after all, an ellipse — is the *second order cone*, to be discussed later.

Proposition 112

Polyhedral and ellipsoidal cones are convex cones.

Proof. Left as an exercise. □

⁴In this class, the issue really only comes up when discussing vector spaces where each element is a matrix, where the norm is the Frobenius norm, and the inner product is a corresponding “Frobenius inner product,” to be defined later. This is relevant in semidefinite programming, for example.

Proposition 113

Let $K \subseteq \mathbb{R}^n$ be a cone. Define

$$K^* \doteq \{\vec{y} \in \mathbb{R}^n \mid \vec{y}^\top \vec{x} \geq 0 \text{ for each } \vec{x} \in K\}. \quad (5.56)$$

Then K^* is a closed convex cone. We call K^* the *dual cone* of K .

Proof. Let $\vec{y}, \vec{z} \in K^*$ and let $\alpha, \beta \geq 0$. Then, for any $\vec{x} \in K$, we have

$$(\alpha\vec{y} + \beta\vec{z})^\top \vec{x} = \underbrace{\alpha}_{\geq 0} \underbrace{(\vec{y}^\top \vec{x})}_{\geq 0} + \underbrace{\beta}_{\geq 0} \underbrace{(\vec{z}^\top \vec{x})}_{\geq 0} \geq 0. \quad (5.57)$$

In particular, this holds for $\beta = 0$ (so that K^* is a cone), and $\alpha \in [0, 1]$ and $\beta = 1 - \alpha$ (so that K^* is convex). Thus, K^* is a convex cone.

Now we want to show that K contains its limits. Let $(\vec{y}_k)_{k=1}^\infty$ be a sequence in K^* that converges to some $\vec{y} \in \mathbb{R}^n$. We want to show that $\vec{y} \in K^*$. Indeed, for any $\vec{x} \in K$, we want to show that $\vec{y}^\top \vec{x} \geq 0$. But this is true because we have

$$\vec{y}^\top \vec{x} = \left(\lim_{k \rightarrow \infty} \vec{y}_k \right)^\top \vec{x} = \lim_{k \rightarrow \infty} \underbrace{\vec{y}_k^\top \vec{x}}_{\geq 0} \geq 0. \quad (5.58)$$

Since \vec{x} was arbitrary, we have $\vec{y} \in K^*$. Thus K^* contains its limits and is a closed cone. \square

A geometric interpretation of the dual cone is that K^* is the intersection of the half-spaces $\mathcal{H}_{\vec{x}} \doteq \{\vec{y} \in \mathbb{R}^n \mid \vec{y}^\top \vec{x} \geq 0\}$ defined by each vector \vec{x} in K .

Below, we provide some examples of cones and their dual cones. The reader is encouraged to verify the following statements.

Example 114.

- (a) The set $\mathbb{R}_+^n \doteq \{\vec{x} \in \mathbb{R}^n \mid x_i \geq 0 \text{ for all } i \in \{1, \dots, n\}\}$ is a convex cone, and its dual cone in \mathbb{R}^n is itself.
- (b) Let $S \doteq \{\vec{x} \in \mathbb{R}^2 \mid x_1 = 0 \text{ or } x_2 = 0\}$. Then S is a cone but is not a convex cone, and the dual cone of S is $\{\vec{0}\}$, the singleton set comprised of the 2-dimensional zero vector.
- (c) Let $S \subseteq \mathbb{R}^n$ be a subspace. Then S is a convex cone, and the orthogonal complement S^\perp of S is the dual cone of S .

Two proper cones with interesting properties that are widely used in convex optimization are the cone of symmetric positive semi-definite matrices and the second-order cone. The propositions below explore their properties.

Proposition 115

Let \mathbb{S}^n be the vector space of $n \times n$ real-valued symmetric matrices equipped with the Frobenius inner product:

$$\langle A, B \rangle_F \doteq \text{tr}(AB) = \sum_{i=1}^n \sum_{j=1}^n A_{ij} B_{ij}, \quad \text{for any } A, B \in \mathbb{S}^n. \quad (5.59)$$

and the Frobenius norm $\|\cdot\|_F$. Let \mathbb{S}_+^n denote the set of all $n \times n$ positive semidefinite matrices.

- (a) \mathbb{S}_+^n is a proper cone in \mathbb{S}^n ;

(b) The dual cone of \mathbb{S}_+^n in \mathbb{S}^n is itself.

To start with, this is an instance of the earlier discussion: not every application of cones will be with reference to \mathbb{R}^n , but could be with reference to another inner product space. Here it is the vector space \mathbb{S}^n with the appropriate inner product and norm. The intuition for why we can do this is that \mathbb{S}^n is a subspace of $\mathbb{R}^{n \times n}$, the space of $n \times n$ matrices. But by stacking up the entries in an $n \times n$ matrix we get an n^2 -dimensional vector, i.e., an element of \mathbb{R}^{n^2} . Indeed, the Frobenius norm and inner product on matrices are exactly the ℓ^2 inner product and norm applied to the “unrolled” matrices in \mathbb{R}^{n^2} . Thus, one can informally view \mathbb{S}^n as a subspace of \mathbb{R}^{n^2} (though remember that it is a vector space in its own right, so that we can define things like interiors and dual cones with respect to it instead of its “parent” space \mathbb{R}^{n^2}), so the same proof techniques and intuitions carry over.

Proof.

(a) To show that \mathbb{S}_+^n is a convex cone, let $A, B \in \mathbb{S}_+^n$ and $\alpha, \beta \geq 0$ be given. We wish to show that $\alpha A + \beta B \in \mathbb{S}_+^n$, which will confirm that \mathbb{S}_+^n is a convex cone. Indeed, $\alpha A + \beta B$ is symmetric as the linear combination of two symmetric matrices. To show that it is positive semidefinite, let $\vec{v} \in \mathbb{R}^n$ be arbitrarily given. Then we have

$$\vec{v}^\top (\alpha A + \beta B) \vec{v} = \underbrace{\alpha}_{\geq 0} \underbrace{\vec{v}^\top A \vec{v}}_{\geq 0} + \underbrace{\beta}_{\geq 0} \underbrace{\vec{v}^\top B \vec{v}}_{\geq 0} \geq 0. \quad (5.60)$$

Since \vec{v} was arbitrarily given, $\alpha A + \beta B$ is positive semidefinite. Thus $\alpha A + \beta B \in \mathbb{S}_+^n$. This holds for $\beta = 0$, so \mathbb{S}_+^n is a cone, and also for $\alpha \in [0, 1]$ and $\beta = 1 - \alpha$, so \mathbb{S}_+^n is convex.

We now show that \mathbb{S}_+^n is pointed, i.e., it contains no lines through the origin. Let $A \in \mathbb{S}_+^n$ be a nonzero matrix. Then $-A \notin \mathbb{S}_+^n$, because there exists $\vec{v} \in \mathbb{R}^n$ such that $\vec{v}^\top A \vec{v} > 0$, at which point $\vec{v}^\top (-A) \vec{v} = -\vec{v}^\top A \vec{v} < 0$, so $-A$ is not positive semidefinite. Thus $-A \notin \mathbb{S}_+^n$, so that for any $A \in \mathbb{S}_+^n$ there exists $\alpha \in \mathbb{R}$ such that $\alpha A \notin \mathbb{S}_+^n$. Thus \mathbb{S}_+^n contains no lines through the origin and is pointed.

We now show that \mathbb{S}_+^n is solid, i.e., has nonempty interior. We show that the open ball in \mathbb{S}^n defined by

$$\mathcal{B} \doteq \left\{ A \in \mathbb{S}^n \mid \|A - I\|_F < \frac{1}{2} \right\} \quad (5.61)$$

is contained in \mathbb{S}_+^n . Indeed, let $A \in \mathcal{B}$. By definition of \mathcal{B} , we have that A is symmetric. Moreover, for each $\vec{v} \in \mathbb{R}^n$ we have

$$\vec{v}^\top A \vec{v} = \vec{v}^\top ((A - I) + I) \vec{v} \quad (5.62)$$

$$= \vec{v}^\top (A - I) \vec{v} + \vec{v}^\top \vec{v} \quad (5.63)$$

$$= \vec{v}^\top (A - I) \vec{v} + \|\vec{v}\|_2^2 \quad (5.64)$$

$$\geq -\|A - I\|_2 \|\vec{v}\|_2^2 + \|\vec{v}\|_2^2 \quad (5.65)$$

$$\geq -\|A - I\|_F \|\vec{v}\|_2^2 + \|\vec{v}\|_2^2 \quad (5.66)$$

$$> -\frac{1}{2} \|\vec{v}\|_2^2 + \|\vec{v}\|_2^2 \quad (5.67)$$

$$= \frac{1}{2} \|\vec{v}\|_2^2. \quad (5.68)$$

For \vec{v} nonzero, we have $\frac{1}{2} \|\vec{v}\|_2^2$, and so $\vec{v}^\top A \vec{v} > 0$. Thus, $A \in \mathbb{S}_+^n$.

Finally, we need to show that \mathbb{S}_+^n is a closed cone. Let $(A_k)_{k=1}^\infty$ be a sequence in \mathbb{S}_+^n that converges to some $A \in \mathbb{S}^n$. We want to show that $A \in \mathbb{S}_+^n$. As the limit of symmetric matrices, A is symmetric. Now for any

$\vec{v} \in \mathbb{R}^n$ we have

$$\vec{v}^\top A \vec{v} = \vec{v}^\top \left(\lim_{k \rightarrow \infty} A_k \right) \vec{v} = \lim_{k \rightarrow \infty} \underbrace{\vec{v}^\top A_k \vec{v}}_{\geq 0} \geq 0. \quad (5.69)$$

Thus $\vec{v}^\top A \vec{v} \geq 0$ for all $\vec{v} \in \mathbb{R}^n$, so that $A \in \mathbb{S}_+^n$. Thus \mathbb{S}_+^n contains its limits and is a closed cone.

We have thus proved that \mathbb{S}_+^n is a convex, pointed, solid, and closed cone, so it is a proper cone.

- (b) We now show that the dual cone of \mathbb{S}_+^n in \mathbb{S}^n is itself. That is, defining the dual cone as $(\mathbb{S}_+^n)^* = \{A \in \mathbb{S}^n \mid \langle A, B \rangle_F \geq 0 \text{ for all } B \in \mathbb{S}_+^n\}$, we want to show that $(\mathbb{S}_+^n)^* = \mathbb{S}_+^n$. To do this, we show that $(\mathbb{S}_+^n)^* \subseteq \mathbb{S}_+^n$ and that $(\mathbb{S}_+^n)^* \supseteq \mathbb{S}_+^n$.

We first show that $(\mathbb{S}_+^n)^* \subseteq \mathbb{S}_+^n$. Fix $A \in (\mathbb{S}_+^n)^*$, and let $\vec{v} \in \mathbb{R}^n$ be given arbitrarily. Then, since $\vec{v}\vec{v}^\top \in \mathbb{S}_+^n$, we have

$$\vec{v}^\top A \vec{v} = \text{tr}(\vec{v}^\top A \vec{v}) \quad (5.70)$$

$$= \text{tr}(A \vec{v} \vec{v}^\top) \quad (5.71)$$

$$= \langle A, \vec{v} \vec{v}^\top \rangle_F \quad (5.72)$$

$$\geq 0, \quad (5.73)$$

where in the first line we use the fact that the trace of a scalar is a scalar, in the second line we use the cyclic trace inequality, and the last inequality is justified because $\vec{v}\vec{v}^\top \in \mathbb{S}_+^n$ and $A \in (\mathbb{S}_+^n)^*$. Since \vec{v} was selected arbitrarily, $\vec{v}^\top A \vec{v} \geq 0$ for all $\vec{v} \in \mathbb{R}^n$. This (along with the fact that A is symmetric) proves that $A \in \mathbb{S}_+^n$. Since A was selected arbitrarily, $(\mathbb{S}_+^n)^* \subseteq \mathbb{S}_+^n$.

Now we show that $\mathbb{S}_+^n \subseteq (\mathbb{S}_+^n)^*$. Let $B \in \mathbb{S}_+^n$. We aim to show that $B \in (\mathbb{S}_+^n)^*$, i.e., that $\langle B, C \rangle_F \geq 0$ for any $C \in \mathbb{S}_+^n$. By the spectral theorem, we may diagonalize $C = \sum_{i=1}^n \lambda_i \vec{v}_i \vec{v}_i^\top$, where $\lambda_i \geq 0$ are the eigenvalues of C and $\vec{v}_i \in \mathbb{R}^n$ are orthonormal eigenvectors of C . Thus we have

$$\langle B, C \rangle_F = \text{tr}(BC) \quad (5.74)$$

$$= \text{tr} \left(B \left(\sum_{i=1}^n \lambda_i \vec{v}_i \vec{v}_i^\top \right) \right) \quad (5.75)$$

$$= \text{tr} \left(\sum_{i=1}^n \lambda_i B \vec{v}_i \vec{v}_i^\top \right) \quad (5.76)$$

$$= \sum_{i=1}^n \lambda_i \text{tr}(B \vec{v}_i \vec{v}_i^\top) \quad (5.77)$$

$$= \sum_{i=1}^n \lambda_i \text{tr}(\vec{v}_i^\top B \vec{v}_i) \quad (5.78)$$

$$= \sum_{i=1}^n \underbrace{\lambda_i}_{\geq 0} \underbrace{\vec{v}_i^\top B \vec{v}_i}_{\geq 0} \quad (5.79)$$

$$\geq 0. \quad (5.80)$$

Thus we have $\langle B, C \rangle_F \geq 0$. Since $C \in \mathbb{S}_+^n$ were arbitrary, we have $B \in (\mathbb{S}_+^n)^*$. Since B were arbitrary, we have $\mathbb{S}_+^n \subseteq (\mathbb{S}_+^n)^*$.

Thus, $(\mathbb{S}_+^n)^* = \mathbb{S}_+^n$.

□

The next example of a cone will be useful when discussing the eponymous *second-order cone programs* (SOCPs).

Definition 116 (Second Order Cone)

The *second-order cone* in \mathbb{R}^{n+1} is the set:

$$K \doteq \{(\vec{x}, t) \in \mathbb{R}^n \times \mathbb{R} \mid \|\vec{x}\|_2 \leq t\}. \quad (5.81)$$

Proposition 117

Let K be the second-order cone in \mathbb{R}^{n+1} .

- (a) K is a proper cone.
- (b) The dual cone of K in \mathbb{R}^{n+1} is itself.

Proof.

- (a) We first show that K is a convex cone. Let $(\vec{x}_1, t_1), (\vec{x}_2, t_2) \in K$ and let $\alpha_1, \alpha_2 \geq 0$. Then

$$\|\alpha_1 \vec{x}_1 + \alpha_2 \vec{x}_2\|_2 \leq \alpha_1 \|\vec{x}_1\|_2 + \alpha_2 \|\vec{x}_2\|_2 \quad (5.82)$$

$$\leq \alpha_1 t_1 + \alpha_2 t_2 \quad (5.83)$$

where the first inequality is by triangle inequality and the second is by definition of a second-order cone. This holds for $\alpha_2 = 0$, showing that K is indeed a cone, and $\alpha_1 \in [0, 1]$ and $\alpha_2 = 1 - \alpha_1$, showing that K is convex. Thus K is a convex cone.

We show that K is pointed, i.e., contains no lines through the origin. Indeed, let $(\vec{x}, t) \in K$ be nonzero. Then either \vec{x} is nonzero or t is nonzero (or both); in the first case, $\|\vec{x}\|_2 > 0$ so since $t \geq \|\vec{x}\|_2$, we have $t > 0$ as well. Thus $t > 0$ in all cases. Thus we certainly *do not* have $\|-\vec{x}\|_2 \leq -t$ (in fact, norms can never be negative) so that $-(\vec{x}, t) = (-\vec{x}, -t) \notin K$. Thus for any $(\vec{x}, t) \in K$ there exists $\alpha \in \mathbb{R}$ such that $\alpha(\vec{x}, t) \notin K$, so K is pointed.

We now show that K is solid, i.e., has nonempty interior. We claim that the open ball in \mathbb{R}^{n+1} of radius 1 centered at $(\vec{0}, 2)$, where $\vec{0}$ is the n -dimensional zero vector, is contained in K . Formally, define

$$\mathcal{B} \doteq \{(\vec{x}, t) \in \mathbb{R}^{n+1} \mid \|(\vec{x}, t) - (\vec{0}, 2)\|_2 < 1\}. \quad (5.84)$$

Let $(\vec{x}, t) \in \mathcal{B}$; we show that $(\vec{x}, t) \in K$. Indeed, we have

$$\|(\vec{x}, t) - (\vec{0}, 2)\|_2 < 1 \quad (5.85)$$

$$\implies \|(\vec{x}, t) - (\vec{0}, 2)\|_2^2 < 1^2 = 1 \quad (5.86)$$

$$\implies \|\vec{x}\|_2^2 + (t - 2)^2 < 1, \quad (5.87)$$

which implies that $\|\vec{x}\|_2^2 < 1$ and $(t - 2)^2 < 1$, namely $t \in (1, 3)$. Thus $\|\vec{x}\|_2^2 < 1 < t$, so $\|\vec{x}\|_2 < 1 < t$, so $\|\vec{x}\|_2 \leq t$, so $(\vec{x}, t) \in K$ as desired. Since $(\vec{x}, t) \in \mathcal{B}$ were arbitrarily chosen, $\mathcal{B} \subseteq K$ and K is solid.

We now show that K is closed, i.e., contains its limits. Let $((\vec{x}_k, t_k))_{k=1}^\infty$ be a sequence in K that converges to some $(\vec{x}, t) \in \mathbb{R}^{n+1}$. We want to show that $(\vec{x}, t) \in K$. Indeed, we have that $(\vec{x}, t) \in K$ if and only if $t - \|\vec{x}\|_2 \geq 0$. We have

$$t - \|\vec{x}\|_2 = \lim_{k \rightarrow \infty} t_k - \left\| \lim_{k \rightarrow \infty} \vec{x}_k \right\|_2 \quad (5.88)$$

$$= \lim_{k \rightarrow \infty} \underbrace{(t_k - \|\vec{x}_k\|_2)}_{\geq 0} \quad (5.89)$$

$$\geq 0. \quad (5.90)$$

Thus $(\vec{x}, t) \in K$, so K contains its limits and is closed.

We have proved that K is a convex, pointed, solid, and closed cone, so it is proper.

- (b) We show that the dual cone of K in \mathbb{R}^{n+1} is K itself. Let $K^* = \{(\vec{y}, s) \in \mathbb{R}^{n+1} \mid (\vec{y}, s)^\top (\vec{x}, t) \geq 0 \text{ for all } (\vec{x}, t) \in K\}$ be the dual cone of K . We first show that $K^* \subseteq K$, then show that $K^* \supseteq K$.

First, to show that $K^* \subseteq K$, fix $(\vec{y}, s) \in K^*$. We want to show that $s \geq \|\vec{y}\|_2$, so that $(\vec{y}, s) \in K^*$. Since $(\vec{0}, 1) \in K$, by definition of K^* we have

$$0 \leq (\vec{y}, s)^\top (\vec{0}, 1) = s. \quad (5.91)$$

Thus, if $\|\vec{y}\|_2 = 0$ (i.e., if $\vec{y} = \vec{0}$), then we have $s \geq \|\vec{y}\|_2$, so that $(\vec{y}, s) \in K$.

Now suppose that $\vec{y} \neq \vec{0}$, so that $\|\vec{y}\|_2 > 0$. Then $(-\vec{y}, \|\vec{y}\|_2) \in K$, so that

$$0 \leq (\vec{y}, s)^\top (-\vec{y}, \|\vec{y}\|_2) = -\|\vec{y}\|_2^2 + s \|\vec{y}\|_2 \quad (5.92)$$

$$\Rightarrow s \|\vec{y}\|_2 \geq \|\vec{y}\|_2^2 \quad (5.93)$$

$$\Rightarrow s \geq \|\vec{y}\|_2. \quad (5.94)$$

Therefore, $(\vec{y}, s) \in K$. Since $(\vec{y}, s) \in K^*$ were arbitrary, we have $K^* \subseteq K$.

Now we want to show that $K \subseteq K^*$. To this end, let $(\vec{y}, s) \in K$. We want to show that $(\vec{y}, s) \in K^*$, or equivalently, $(\vec{x}, t)^\top (\vec{y}, s) \geq 0$ for all $(\vec{x}, t) \in K$. Indeed,

$$(\vec{x}, t)^\top (\vec{y}, s) = \vec{x}^\top \vec{y} + st \geq -\|\vec{x}\|_2 \|\vec{y}\|_2 - st \geq 0 \quad (5.95)$$

where the first inequality follows by Cauchy-Schwarz inequality, and the second inequality follows from the fact that since $(\vec{x}, t), (\vec{y}, s) \in K$ we have $\|\vec{x}\|_2 \leq t$ and $\|\vec{y}\|_2 \leq s$. This shows that $(\vec{y}, s) \in K^*$, so that $K \subseteq K^*$.

We have shown that $K \supseteq K^*$ and $K \subseteq K^*$, so $K = K^*$.

□

Theorem 118

Let $K \subseteq \mathbb{R}^n$ be a non-empty closed convex cone. Then $(K^*)^* = K$.

Proof. We want to show that $K \subseteq (K^*)^*$ and $K \supseteq (K^*)^*$.

First, we want to show that $K \subseteq (K^*)^*$. Fix $\vec{x} \in K$. We want to show that $\vec{x} \in (K^*)^*$, which means that $\vec{x}^\top \vec{y} \geq 0$ for any $\vec{y} \in K^*$. For this \vec{y} , we have $\vec{y}^\top \vec{z} \geq 0$ for all $\vec{z} \in K$. But this includes $\vec{z} = \vec{x}$, so $\vec{x}^\top \vec{y} = \vec{y}^\top \vec{x} \geq 0$. Thus $\vec{x} \in (K^*)^*$. We conclude that $K \subseteq (K^*)^*$. (Note that $K \subseteq (K^*)^*$ actually holds for any cone K , not just non-empty closed convex cones.)

Next, we want to show that $(K^*)^* \subseteq K$. Suppose for the sake of contradiction that $(K^*)^* \not\subseteq K$. Then there exists $\vec{y} \in (K^*)^*$ such that $\vec{y} \notin K$. Since $\vec{y} \notin K$, we have $\vec{y} \neq \vec{0}$. Because K is a closed convex cone, it is a closed convex set. Since $\{\vec{y}\}$ and K are two disjoint closed convex sets, the Separating Hyperplane Theorem tells us that there exists some

nonzero $\vec{w} \in \mathbb{R}^n$ and $c \in \mathbb{R}$ such that $\vec{w}^\top \vec{x} > c$ for all $\vec{x} \in K$, and $\vec{w}^\top \vec{y} < c$. Since $\vec{0} \in K$, we have $0 = \vec{w}^\top \vec{0} > c$, i.e., $c < 0$, so $\vec{w}^\top \vec{y} < c < 0$. Since K is a cone, for any $\alpha > 0$ we have

$$\vec{w}^\top (\alpha \vec{x}) > c \quad (5.96)$$

$$\Rightarrow \alpha \vec{w}^\top \vec{x} > c \quad (5.97)$$

$$\Rightarrow \vec{w}^\top \vec{x} > \frac{c}{\alpha}. \quad (5.98)$$

By taking $\alpha \rightarrow \infty$ we get that $\vec{w}^\top \vec{x} \geq 0$ for any $\vec{x} \in K$. Thus $\vec{w} \in K^*$. But we must have $\vec{w}^\top \vec{y} < c < 0$, so $\vec{y} \notin (K^*)^*$, a contradiction. Thus $(K^*)^* \subseteq K$ and so $(K^*)^* = K$. \square

The above content is optional/out of scope for this semester, but now we resume the required/in scope content.

5.2 Convex Functions

In this section, we define convex and concave functions, and introduce their properties. At the end, we also deliberate on the special and important example of *affine functions*.

Definition 119 (Convex and Concave Functions)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function. We say that f is *convex* if the domain of f , say Ω , is convex, and, for any $\vec{x}_1, \vec{x}_2 \in \Omega$ and $\theta \in [0, 1]$, we have

$$f(\theta \vec{x}_1 + (1 - \theta) \vec{x}_2) \leq \theta f(\vec{x}_1) + (1 - \theta) f(\vec{x}_2). \quad (5.99)$$

A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is *concave* if $-f$ is convex.


Equation (5.99) is also called *Jensen's inequality* and is equivalent to the following, seemingly more general statement.

Theorem 120 (Jensen's Inequality)

Let Ω be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a convex function. Let $\vec{x}_1, \dots, \vec{x}_k \in \Omega$, and let $\theta_1, \dots, \theta_k \in [0, 1]$ such that $\sum_{i=1}^k \theta_i = 1$. Then

$$f\left(\sum_{i=1}^k \theta_i \vec{x}_i\right) \leq \sum_{i=1}^k \theta_i f(\vec{x}_i). \quad (5.100)$$

We can think about this result in terms of a picture.



The prototypical convex function has a “bowl-shaped” graph, and taking a weighted average of two (or any finite number) of points will mean we land in the bowl. In particular, taking a weighted average of any number of function values $f(\vec{x}_1), \dots, f(\vec{x}_k)$ will always give a larger number than applying the function f to the same weighted averages of the points $\vec{x}_1, \dots, \vec{x}_k$. Put more simply, if f is convex then the chord joining the points $(\vec{x}_1, f(\vec{x}_1))$ and $(\vec{x}_2, f(\vec{x}_2))$ always lies *above* the graph of f . Similarly, if f is concave then the chord joining the points $(\vec{x}_1, f(\vec{x}_1))$ and $(\vec{x}_2, f(\vec{x}_2))$ always lies *below* the graph of f .

From the picture, it may be intuitively clear that it is hard to construct convex functions f with multiple global minima. We will come back to this idea later.


It may be useful to connect the notion of convex function and convex set. For this, we will define the epigraph.

Definition 121 (Epigraph)

Let Ω be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a convex function. The *epigraph* of f , denoted $\text{epi}(f) \subseteq \Omega \times \mathbb{R}$, is defined as

$$\text{epi}(f) = \{(\vec{x}, t) \mid \vec{x} \in \Omega, t \geq f(\vec{x})\}. \quad (5.101)$$

Geometrically, the epigraph is all points in $\Omega \times \mathbb{R}$ that lie above the graph of f .

**Proposition 122**

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a function. Then f is a convex function if and only if $\text{epi}(f)$ is a convex set.


Proof. Left as an exercise. □

Theorem 123 (First-Order Condition for Convexity)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a differentiable function. Then f is convex if and only if for all $\vec{x}, \vec{y} \in \Omega$, we have

$$f(\vec{y}) \geq f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x}). \quad (5.102)$$

Note that this latter term is the first-order Taylor expansion of f around \vec{x} evaluated at \vec{y} , i.e., $\hat{f}_1(\vec{y}; \vec{x}) = f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x})$. The graph of $\hat{f}_1(\cdot; \vec{x})$ is the tangent line to the graph of f at the point $(\vec{x}, f(\vec{x}))$. So another characterization of convex functions is that their graphs lie above their tangent lines.



Proof. First suppose f is convex. Then for any $h \in (0, 1)$, we have

$$f(h\bar{y} + (1-h)\bar{x}) \leq hf(\bar{y}) + (1-h)f(\bar{x}) \quad (5.103)$$

$$= hf(\bar{y}) + f(\bar{x}) - hf(\bar{x}) \quad (5.104)$$

$$= f(\bar{x}) + h(f(\bar{y}) - f(\bar{x})) \quad (5.105)$$

$$\implies f(h\bar{y} + (1-h)\bar{x}) - f(\bar{x}) \leq h(f(\bar{y}) - f(\bar{x})) \quad (5.106)$$

$$\implies \frac{f(h\bar{y} + (1-h)\bar{x}) - f(\bar{x})}{h} \leq f(\bar{y}) - f(\bar{x}) \quad (5.107)$$

$$\implies f(\bar{y}) \geq f(\bar{x}) + \frac{f(h\bar{y} + (1-h)\bar{x}) - f(\bar{x})}{h} \quad (5.108)$$

$$= f(\bar{x}) + \frac{f(\bar{x} + h(\bar{y} - \bar{x})) - f(\bar{x})}{h}. \quad (5.109)$$

To summarize, for any $h \in (0, 1)$ we have that

$$f(\bar{y}) \geq f(\bar{x}) + \frac{f(\bar{x} + h(\bar{y} - \bar{x})) - f(\bar{x})}{h}. \quad (5.110)$$

Taking the limit $h \rightarrow 0$ on both sides, we get

$$f(\bar{y}) \geq \lim_{h \rightarrow 0} \left\{ f(\bar{x}) + \frac{f(\bar{x} + h(\bar{y} - \bar{x})) - f(\bar{x})}{h} \right\} \quad (5.111)$$

$$= f(\bar{x}) + \lim_{h \rightarrow 0} \frac{f(\bar{x} + h(\bar{y} - \bar{x})) - f(\bar{x})}{h} \quad (5.112)$$

$$= f(\bar{x}) + [\nabla f(\bar{x})]^\top (\bar{y} - \bar{x}) \quad (5.113)$$

as desired. Here the last equality is because the limit is interpreted as a directional derivative, and it has already been shown that directional derivatives are equal to inner products of the gradient with the direction vector.

For the other direction, let $\theta \in [0, 1]$ and let $\bar{z} = \theta\bar{x} + (1-\theta)\bar{y}$. We have

$$f(\bar{x}) \geq f(\bar{z}) + [\nabla f(\bar{z})]^\top (\bar{x} - \bar{z}) \quad (5.114)$$

$$f(\vec{y}) \geq f(\vec{z}) + [\nabla f(\vec{z})]^\top (\vec{y} - \vec{z}). \quad (5.115)$$

Adding θ times the first equation to $(1 - \theta)$ times the second equation, we get

$$\theta f(\vec{x}) + (1 - \theta)f(\vec{y}) \geq \theta f(\vec{z}) + (1 - \theta)f(\vec{z}) + \theta[\nabla f(\vec{z})]^\top (\vec{x} - \vec{z}) + (1 - \theta)[\nabla f(\vec{z})]^\top (\vec{y} - \vec{z}) \quad (5.116)$$

$$= f(\vec{z}) + [\nabla f(\vec{z})]^\top (\theta\vec{x} + (1 - \theta)\vec{y} - \vec{z}) \quad (5.117)$$

$$= f(\vec{z}) + [\nabla f(\vec{z})]^\top (\vec{z} - \vec{z}) \quad (5.118)$$

$$= f(\vec{z}) + [\nabla f(\vec{z})]^\top \vec{0} \quad (5.119)$$

$$= f(\vec{z}) \quad (5.120)$$

$$= f(\theta\vec{x} + (1 - \theta)\vec{y}). \quad (5.121)$$

□

We also have a corresponding second-order condition.

Theorem 124 (Second-Order Condition for Convexity)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a twice-differentiable function. Then f is convex if and only if for all $\vec{x} \in \Omega$, we have

$$\nabla^2 f(\vec{x}) \succeq 0, \quad (5.122)$$

i.e., $\nabla^2 f(\vec{x})$ is PSD for each $\vec{x} \in \Omega$.

Corollary 125. Let $Q \in \mathbb{S}^n$ be a symmetric matrix, let $\vec{b} \in \mathbb{R}^n$, and let $c \in \mathbb{R}$. The quadratic form

$$f(\vec{x}) = \vec{x}^\top Q \vec{x} + \vec{b}^\top \vec{x} + c \quad (5.123)$$

is convex if and only if Q is PSD.

Lastly, we identify a strengthened condition of convexity which allows for stronger guarantees later down the line.

Definition 126 (Strictly Convex Function)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function. We say that f is *strictly convex* if the domain of f , say Ω , is convex, and, for any $\vec{x}_1 \neq \vec{x}_2 \in \Omega$ and $\theta \in (0, 1)$, we have

$$f(\theta\vec{x}_1 + (1 - \theta)\vec{x}_2) < \theta f(\vec{x}_1) + (1 - \theta)f(\vec{x}_2). \quad (5.124)$$

A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is *strictly concave* if $-f$ is strictly convex.

And correspondingly, we have the first-order and second-order conditions.

Theorem 127 (First-Order Condition for Strict Convexity)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a differentiable function. Then f is strictly convex if and only if for all $\vec{x} \neq \vec{y} \in \Omega$, we have

$$f(\vec{y}) > f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x}). \quad (5.125)$$

Theorem 128 (Second-Order Condition for Strict Convexity)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a twice-differentiable function such that $\nabla^2 f(\vec{x}) \succ 0$ for all $\vec{x} \in \Omega$, i.e., $\nabla^2 f(\vec{x})$ is PD for all $\vec{x} \in \Omega$. Then f is strictly convex.

Notice that this is *not* an if-and-only-if! As an example, take the scalar function $f(x) = x^4$. Then $f''(0) = 0$, so it is not true that $f''(x) \succ 0$ for all $x \in \Omega = \mathbb{R}$, but f is strictly convex.

5.2.1 Affine Functions

Finally, we spend some time on the special case of *affine* functions, which are the only functions that are both convex and concave. These are of special importance in the remainder of the course, being used or referenced in almost every topic henceforth.

Definition 129 (Affine Functions)

- (a) A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be *affine* if there exists some vector $\vec{a} \in \mathbb{R}^n$ and some scalar $b \in \mathbb{R}$ such that for any $\vec{x} \in \mathbb{R}^n$, we have

$$f(\vec{x}) = \vec{a}^\top \vec{x} + b. \quad (5.126)$$

- (b) A function $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is said to be *affine* if there exists some matrix $A \in \mathbb{R}^{m \times n}$ and some vector $\vec{b} \in \mathbb{R}^m$ such that for any $\vec{x} \in \mathbb{R}^n$, we have

$$\vec{f}(\vec{x}) = A\vec{x} + \vec{b}. \quad (5.127)$$

- (c) A function $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is said to be *affine* if there exists some matrix $A \in \mathbb{R}^{m \times n}$ and scalar $b \in \mathbb{R}$ such that for any $X \in \mathbb{R}^{m \times n}$, we have

$$f(X) = \sum_{i=1}^m \sum_{j=1}^n A_{ij} X_{ij} + b = \text{tr}(A^\top X) + b. \quad (5.128)$$

Note that a given function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is affine if and only if the function $g: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by $g(\vec{x}) = f(\vec{x}) - f(\vec{0})$ is linear. An analogous result holds for other types of affine functions.

Below, we show that a scalar-valued affine function is one that is both convex and concave, while a vector-valued affine function is one whose component functions are all both convex and concave. Analogous results hold for affine functions whose inputs and outputs are both matrices.

Proposition 130

- (a) A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is affine if and only if it is both convex and concave, i.e., for any $\alpha \in [0, 1]$ and $\vec{x}, \vec{y} \in \mathbb{R}^n$, we have

$$f(\alpha\vec{x} + (1-\alpha)\vec{y}) = \alpha f(\vec{x}) + (1-\alpha)f(\vec{y}). \quad (5.129)$$

- (b) A function $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is affine if and only if each component function of f is both convex and concave, i.e., for any $\alpha \in [0, 1]$ and $\vec{x}, \vec{y} \in \mathbb{R}^n$, we have

$$\vec{f}(\alpha\vec{x} + (1-\alpha)\vec{y}) = \alpha\vec{f}(\vec{x}) + (1-\alpha)\vec{f}(\vec{y}). \quad (5.130)$$

(c) A function $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is affine if and only if it is both convex and concave, i.e., for any $\alpha \in [0, 1]$ and $X, Y \in \mathbb{R}^{m \times n}$, we have

$$f(\alpha X + (1 - \alpha)Y) = \alpha f(X) + (1 - \alpha)f(Y). \quad (5.131)$$

Proof. We prove (a); the claims (b) and (c) follow similarly.

Suppose first that f is affine. Then there exists $\vec{a} \in \mathbb{R}^n$ and $b \in \mathbb{R}$ such that for each $\vec{x}, \vec{y} \in \mathbb{R}^n$ and $\alpha \in [0, 1]$, we have

$$f(\alpha \vec{x} + (1 - \alpha)\vec{y}) = \vec{a}^\top (\alpha \vec{x} + (1 - \alpha)\vec{y}) + b \quad (5.132)$$

$$= \alpha \vec{a}^\top \vec{x} + (1 - \alpha)\vec{a}^\top \vec{y} + b \quad (5.133)$$

$$= \alpha(\vec{a}^\top \vec{x} + b) + (1 - \alpha)(\vec{a}^\top \vec{y} + b) \quad (5.134)$$

$$= \alpha f(\vec{x}) + (1 - \alpha)f(\vec{y}). \quad (5.135)$$

Conversely, suppose that

$$f(\alpha \vec{x} + (1 - \alpha)\vec{y}) = \alpha f(\vec{x}) + (1 - \alpha)f(\vec{y}), \quad \text{for all } \alpha \in [0, 1] \text{ and } \vec{x}, \vec{y} \in \mathbb{R}^n. \quad (5.136)$$

To show that f is affine, it suffices to show that $g: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by $g(\vec{x}) = f(\vec{x}) - f(\vec{0})$ is linear (and thus can be written as an inner product against a vector \vec{a}). We first show that $g(r\vec{x}) = rg(\vec{x})$ for any $r \in \mathbb{R}$ and $\vec{x} \in \mathbb{R}^n$. We break this problem up into three cases, each building on the other.

Case 1. Suppose that $r \in [0, 1]$. Then $r\vec{x}$ can be expressed as a convex combination of $\vec{0}$ and \vec{x} : that is,

$$r\vec{x} = \alpha \vec{x} + (1 - \alpha)\vec{0}, \quad \text{where } \alpha = r \in (0, 1). \quad (5.137)$$

(Yes, this is a simpler step, but we build on it in the later parts.) With this, we have

$$f(r\vec{x}) = f(\alpha \vec{x} + (1 - \alpha)\vec{0}) = \alpha f(\vec{x}) + (1 - \alpha)f(\vec{0}). \quad (5.138)$$

Thus, we obtain

$$g(r\vec{x}) = f(r\vec{x}) - f(\vec{0}) \quad (5.139)$$

$$= \alpha f(\vec{x}) + (1 - \alpha)f(\vec{0}) - f(\vec{0}) \quad (5.140)$$

$$= \alpha f(\vec{x}) - \alpha f(\vec{0}) \quad (5.141)$$

$$= \alpha(f(\vec{x}) - f(\vec{0})) \quad (5.142)$$

$$= \alpha g(\vec{x}). \quad (5.143)$$

Case 2. Now suppose that $r \in (1, \infty)$. Then \vec{x} can be expressed as a convex combination of $\vec{0}$ and $r\vec{x}$: that is,

$$\vec{x} = \alpha(r\vec{x}) + (1 - \alpha)\vec{0}, \quad \text{where } \alpha = \frac{1}{r} \in (0, 1). \quad (5.144)$$

Thus we have

$$f(\vec{x}) = f\left(\frac{1}{r} \cdot r\vec{x}\right) = f\left(\frac{1}{r} \cdot r\vec{x} + \left(1 - \frac{1}{r}\right)\vec{0}\right) = \frac{1}{r}f(r\vec{x}) + \left(1 - \frac{1}{r}\right)f(\vec{0}). \quad (5.145)$$

Multiplying both sides by r and plugging it into the previous calculation, we get

$$g(r\vec{x}) = f(r\vec{x}) - f(\vec{0}) \quad (5.146)$$

$$= rf(\vec{x}) - r \left(1 - \frac{1}{r}\right) f(\vec{0}) - f(\vec{0}) \quad (5.147)$$

$$= rf(\vec{x}) - (r-1)f(\vec{0}) - f(\vec{0}) \quad (5.148)$$

$$= rf(\vec{x}) - rf(\vec{0}) \quad (5.149)$$

$$= r(f(\vec{x}) - f(\vec{0})) \quad (5.150)$$

$$= rg(\vec{x}). \quad (5.151)$$

Case 3. Now suppose that $r \in (-\infty, 0)$. Then $\vec{0}$ can be expressed as a convex combination of \vec{x} and $r\vec{x}$: that is,

$$0 = \alpha(r\vec{x}) + (1-\alpha)\vec{x}, \text{ where } \alpha = \frac{1}{1-r} \in (0, 1). \quad (5.152)$$

Thus we have

$$f(\vec{0}) = f\left(\frac{1}{1-r}(r\vec{x}) + \left(1 - \frac{1}{1-r}\right)\vec{x}\right) = \frac{1}{1-r}f(r\vec{x}) + \frac{1-r-1}{1-r}f(\vec{x}) = \frac{1}{1-r}f(r\vec{x}) - \frac{r}{1-r}f(\vec{x}). \quad (5.153)$$

Multiplying both sides by $1-r$ and plugging it into the previous calculation, we get

$$g(r\vec{x}) = f(r\vec{x}) - f(\vec{0}) \quad (5.154)$$

$$= rf(\vec{x}) + (1-r)f(\vec{0}) - f(\vec{0}) \quad (5.155)$$

$$= rf(\vec{x}) - rf(\vec{0}) \quad (5.156)$$

$$= r(f(\vec{x}) - f(\vec{0})) \quad (5.157)$$

$$= rg(\vec{x}). \quad (5.158)$$

This proves that, no matter the value of r , we have $g(r\vec{x}) = rg(\vec{x})$.

Now we show that $g(\vec{x} + \vec{y}) = g(\vec{x}) + g(\vec{y})$. With the preceding result in hand, this proof is much shorter:

$$g(\vec{x} + \vec{y}) = 2g\left(\frac{1}{2}\vec{x} + \frac{1}{2}\vec{y}\right) \quad (5.159)$$

$$= 2\left[f\left(\frac{1}{2}\vec{x} + \frac{1}{2}\vec{y}\right) - f(\vec{0})\right] \quad (5.160)$$

$$= 2\left[\frac{1}{2}f(\vec{x}) + \frac{1}{2}f(\vec{y}) - f(\vec{0})\right] \quad (5.161)$$

$$= 2\left[\frac{1}{2}(f(\vec{x}) - f(\vec{0})) + \frac{1}{2}(f(\vec{y}) - f(\vec{0}))\right] \quad (5.162)$$

$$= 2\left[\frac{1}{2}g(\vec{x}) + \frac{1}{2}g(\vec{y})\right] \quad (5.163)$$

$$= g(\vec{x}) + g(\vec{y}). \quad (5.164)$$

Thus we proved that g is linear, so f is affine. This is a full proof of (a), and (b) and (c) can be proved in almost exactly the same way. \square

5.3 Convex Optimization Problems

This section will lay out some of the key properties of the main class of optimization problems we are interested in — convex optimization problems.

Definition 131 (Convex Optimization Problem)

Let $\Omega \subseteq \mathbb{R}^n$ be a set and let $f: \Omega \rightarrow \mathbb{R}$ be a function. We say that the problem

$$\min_{\vec{x} \in \Omega} f(\vec{x}) \quad (5.165)$$

is a *convex optimization problem* if Ω is a convex set and f is a convex function.

Note that this applies to other kinds of constraint sets too — in particular, those of the “standard” form “ $f_i(\vec{x}) \leq 0$ for all i and $h_j(\vec{x}) = 0$ for all j ” still define a feasible region Ω and thus can furnish a convex optimization problem. In particular, we have the following result.

Theorem 132 (When is Standard Form Optimization Problem Convex?)

Let $f_1, \dots, f_m, h_1, \dots, h_p: \mathbb{R}^n \rightarrow \mathbb{R}$ be functions. The feasible set

$$\Omega = \left\{ \vec{x} \in \mathbb{R}^n \mid \begin{array}{l} f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\} \end{array} \right\} \quad (5.166)$$

is convex if each f_i is convex and each h_j is affine. Consequently, the problem

$$\begin{aligned} \min_{\vec{x} \in \mathbb{R}^n} \quad & f_0(\vec{x}) \\ \text{s.t.} \quad & f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}. \end{aligned} \quad (5.167)$$

is a convex optimization problem if f_0, f_1, \dots, f_m are convex functions and h_1, \dots, h_p are affine functions.

Proof. Left as exercise. □

Now, we can establish the first-order condition for optimality within a convex problem. This is one of the main theorems of convex analysis.

Theorem 133 (First-Order Conditions for Optimality in Convex Problem)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a differentiable convex function. Let $\vec{x}^* \in \Omega$ be such that $\nabla f(\vec{x}^*) = \vec{0}$. Then $\vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \Omega} f(\vec{x})$, i.e., \vec{x}^* is a global minimizer of f .

Proof. Let \vec{y} be any other point in Ω . Then

$$f(\vec{y}) \geq f(\vec{x}^*) + [\nabla f(\vec{x}^*)]^\top (\vec{y} - \vec{x}^*) \quad (5.168)$$

$$= f(\vec{x}^*) + \vec{0}^\top (\vec{y} - \vec{x}^*) \quad (5.169)$$

$$= f(\vec{x}^*). \quad (5.170)$$

This proves that $\vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \Omega} f(\vec{x})$ as desired. □

A generalization of this statement is that all local minimizers are global minimizers.

Theorem 134 (For Convex Functions, Local Minima are Global Minima)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a convex function. Let $\bar{x}^* \in \Omega$ be such that there exists some $\epsilon > 0$ such that if $\bar{x} \in \Omega$ has $\|\bar{x} - \bar{x}^*\|_2 \leq \epsilon$ then $f(\bar{x}^*) \leq f(\bar{x})$. Then \bar{x}^* is a minimizer of f over Ω .

Proof. Left as exercise. □

When is the global minimizer unique? We can justify this using strict convexity.

Theorem 135 (Strictly Convex Functions have Unique Minimizers)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a strictly convex function. Then f has *at most* one global minimizer.

Proof. Left as exercise. □

For an example of a strictly convex function with one global minimizer, take $f(x) = x^4$, which is minimized at $x = 0$. For an example of a strictly convex function with no global minimizers, take $f(x) = e^x$.

5.4 Solving Convex Optimization Problems

To construct a first attempt at systematically solving convex optimization problems, we first need to define active and inactive constraints.

Definition 136 (Types of Constraints)

Consider a problem of the form

$$\begin{aligned} \min_{\bar{x} \in \mathbb{R}^n} \quad & f_0(\bar{x}) & (5.171) \\ \text{s.t.} \quad & f_i(\bar{x}) \leq 0, & \forall i \in \{1, \dots, m\} \\ & h_j(\bar{x}) = 0, & \forall j \in \{1, \dots, p\}. \end{aligned}$$

Fix a feasible $\bar{x}_0 \in \mathbb{R}^n$. The *inequality* constraint $f_k(\bar{x}) \leq 0$ is *active* at \bar{x}_0 if $f_k(\bar{x}_0) = 0$, and *inactive* at \bar{x}_0 otherwise, i.e., $f_k(\bar{x}_0) < 0$.

We can use this to formulate a strategy for solving convex optimization problems. Recall that for a convex problem $\min_{\bar{x} \in \Omega} f_0(\bar{x})$ which has a solution \bar{x}^* , either $\nabla f(\bar{x}^*) = \vec{0}$ or \bar{x}^* is on the boundary of Ω . The boundary of Ω is any point in which any inequality constraint is active. This allows us to systematically find solutions to convex optimization problems.

Problem Solving Strategy 137. *To solve a convex optimization program, we can do the following.*

1. Iterate through all 2^m subsets $S \subseteq \{1, \dots, m\}$ of constraints which might be active at optimum.
2. For each S :
 - (i) Solve the modified problem

$$\min_{\bar{x} \in \mathbb{R}^n} f_0(\bar{x}) \tag{5.172}$$

$$\begin{aligned} \text{s.t. } f_i(\vec{x}) &= 0, & \forall i \in S \\ h_j(\vec{x}) &= 0, & \forall j \in \{1, \dots, p\}, \end{aligned}$$

i.e., solve the problem where you pretend that all inequality constraints in S are met with equality and pretend that the other inequality constraints don't exist. This gives some solutions \vec{x}_S^ .*

- (ii) *If there is a solution \vec{x}_S^* which is feasible for the original problem, write down the value of $f_0(\vec{x}_S^*)$. Otherwise, ignore it.*
- (iii) *After iterating through all \vec{x}_S^* which are feasible for the original problem, take the one(s) with the best objective value $f_0(\vec{x}_S^*)$ as the optimal solution(s) to the original problem.*

Predictably, this problem solving strategy is exponentially hard as the number of inequality constraints increases. Even if solving the “inner” equality-constrained minimization problems is easy (as it often is), the whole procedure is untenable for large-scale problems. In future chapters, we will develop better analytic and algorithmic ways to solve convex optimization problems.

5.5 Problem Transformations and Reparameterizations

In this section, we discuss various transformations, or reparameterizations, of generic optimization problems, which allow us to go between many equivalent formulations of a problem. Some of them will even allow us to turn non-convex problems into convex problems.

5.5.1 Monotone Transformations of the Objective Function

The core of this technique is the following result.

Proposition 138 (Positive Monotone Transformations Don't Affect Optimizers)

Let Ω be any set, let $f_0: \Omega \rightarrow \mathbb{R}$ be any function. Define $f_0(\Omega) \doteq \{f_0(\vec{x}) \mid \vec{x} \in \Omega\} \subseteq \mathbb{R}$.

- (a) Let $\phi: f_0(\Omega) \rightarrow \mathbb{R}$ be any monotonically increasing function. Then

$$\operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x}) = \operatorname{argmin}_{\vec{x} \in \Omega} \phi(f_0(\vec{x})) \quad \text{and} \quad \operatorname{argmax}_{\vec{x} \in \Omega} f_0(\vec{x}) = \operatorname{argmax}_{\vec{x} \in \Omega} \phi(f_0(\vec{x})). \quad (5.173)$$

- (b) Let $\psi: f_0(\Omega) \rightarrow \mathbb{R}$ be any monotonically decreasing function. Then

$$\operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x}) = \operatorname{argmax}_{\vec{x} \in \Omega} \psi(f_0(\vec{x})) \quad \text{and} \quad \operatorname{argmax}_{\vec{x} \in \Omega} f_0(\vec{x}) = \operatorname{argmin}_{\vec{x} \in \Omega} \psi(f_0(\vec{x})). \quad (5.174)$$

Proof. We only prove the very first equality; the rest follow similarly. We have

$$\vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \Omega} f_0(\vec{x}) \quad (5.175)$$

$$\iff f_0(\vec{x}^*) \leq f_0(\vec{x}), \quad \forall \vec{x} \in \Omega \quad (5.176)$$

$$\iff \phi(f_0(\vec{x}^*)) \leq \phi(f_0(\vec{x})), \quad \forall \vec{x} \in \Omega \quad (5.177)$$

$$\iff \vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \Omega} \phi(f_0(\vec{x})). \quad (5.178)$$

□

The first equality will be by far the most important for us, though the others might also be situationally useful. This proposition is why doing things like squaring the norm in least squares won't affect the solutions we get, i.e.,

$$\operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{b}\|_2 = \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{b}\|_2^2. \quad (5.179)$$

But the fact that ϕ is monotonic *when restricted to* $f_0(\Omega)$ is quite crucial; indeed, we have that

$$\operatorname{argmin}_{x \in \mathbb{R}} x = \emptyset \quad \text{but} \quad \operatorname{argmin}_{x \in \mathbb{R}} x^2 = \{0\}. \quad (5.180)$$

This is because the function $u \mapsto u^2$ is not monotonic in general, although it is monotonic on the non-negative real numbers \mathbb{R}_+ . It just so happens that $\|\cdot\|_2$ only outputs non-negative numbers, so actually in the case of least squares we have $f_0(\Omega) = \mathbb{R}_+$ and the proposition applies.

Example 139 (Logistic Regression). A more non-trivial example is logistic regression. First, define $\sigma: \mathbb{R} \rightarrow (0, 1)$ by

$$\sigma(x) \doteq \frac{1}{1 + e^{-x}}. \quad (5.181)$$

Suppose we have data points $\vec{x}_1, \dots, \vec{x}_n \in \mathbb{R}^d$ and accompanying labels $y_1, \dots, y_n \in \{0, +1\}$. Suppose that the conditional probability that $y_i = 1$ given \vec{x}_i is given by

$$\mathbb{P}_{\vec{w}_0}[y_i = 1 \mid \vec{x}_i] = \sigma(\vec{x}_i^\top \vec{w}_0) \quad (5.182)$$

for some $\vec{w}_0 \in \mathbb{R}^d$. We wish to recover \vec{w}_0 , and thus recover the generative model $\mathbb{P}_{\vec{w}_0}[y \mid \vec{x}]$. We do this by maximum likelihood estimation. Writing out the problem, we get

$$\vec{w}^* \in \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \mathbb{P}_{\vec{w}}[y_1, \dots, y_n \mid \vec{x}_1, \dots, \vec{x}_n] \quad (5.183)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \prod_{i=1}^n \mathbb{P}_{\vec{w}}[y_i \mid \vec{x}_i] \quad (5.184)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \left(\prod_{\substack{i=1 \\ y_i=0}}^n \mathbb{P}_{\vec{w}}[y_i = 0 \mid \vec{x}_i] \right) \left(\prod_{\substack{i=1 \\ y_i=1}}^n \mathbb{P}_{\vec{w}}[y_i = 1 \mid \vec{x}_i] \right) \quad (5.185)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \prod_{i=1}^n \mathbb{P}_{\vec{w}}[y_i = 1 \mid \vec{x}_i]^{y_i} \mathbb{P}_{\vec{w}}[y_i = 0 \mid \vec{x}_i]^{1-y_i} \quad (5.186)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \prod_{i=1}^n \sigma(\vec{x}_i^\top \vec{w})^{y_i} (1 - \sigma(\vec{x}_i^\top \vec{w}))^{1-y_i}. \quad (5.187)$$

Now, the σ function is *very* non-convex. The product of σ s is *also* very non-convex. Thus, it seems intractable to solve this problem. But note that the objective function takes values in $(0, 1)$. We use the above proposition with the function $x \mapsto \log(x)$, which is monotonically increasing on $(0, 1)$. We obtain

$$\operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \prod_{i=1}^n \sigma(\vec{x}_i^\top \vec{w})^{y_i} (1 - \sigma(\vec{x}_i^\top \vec{w}))^{1-y_i} \quad (5.188)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \log \left(\prod_{i=1}^n \sigma(\vec{x}_i^\top \vec{w})^{y_i} (1 - \sigma(\vec{x}_i^\top \vec{w}))^{1-y_i} \right) \quad (5.189)$$

$$= \operatorname{argmax}_{\vec{w} \in \mathbb{R}^d} \sum_{i=1}^n (y_i \log(\sigma(\vec{x}_i^\top \vec{w})) + (1 - y_i) \log(1 - \sigma(\vec{x}_i^\top \vec{w}))) \quad (5.190)$$

$$= \operatorname{argmin}_{\vec{w} \in \mathbb{R}^d} \left\{ - \sum_{i=1}^n (y_i \log(\sigma(\vec{x}_i^\top \vec{w})) + (1 - y_i) \log(1 - \sigma(\vec{x}_i^\top \vec{w}))) \right\} \quad (5.191)$$

$$= \operatorname{argmin}_{\vec{w} \in \mathbb{R}^d} \sum_{i=1}^n (-y_i \log(\sigma(\vec{x}_i^\top \vec{w})) - (1 - y_i) \log(1 - \sigma(\vec{x}_i^\top \vec{w}))). \quad (5.192)$$

In the penultimate line we used another one of the equalities in the proposition with the monotonically decreasing function $\psi(x) = -x$. Thus, logistic regression reduces to minimizing the objective function

$$f_0(\vec{w}) = \sum_{i=1}^n (-y_i \log(\sigma(\vec{x}_i^\top \vec{w})) - (1 - y_i) \log(1 - \sigma(\vec{x}_i^\top \vec{w}))). \quad (5.193)$$

This is the so-called *cross-entropy* loss.

Computing the gradient and Hessian of this function is an exercise, but the result is

$$\nabla f_0(\vec{w}) = \sum_{i=1}^n \vec{x}_i (\sigma(\vec{x}_i^\top \vec{w}) - y_i) \quad (5.194)$$

$$\nabla^2 f_0(\vec{w}) = \sum_{i=1}^n \sigma(\vec{x}_i^\top \vec{w})(1 - \sigma(\vec{x}_i^\top \vec{w})) \vec{x}_i \vec{x}_i^\top. \quad (5.195)$$

Because $\sigma(x) \in (0, 1)$, the above Hessian is a non-negative weighted sum of positive semidefinite matrices $\vec{x}_i \vec{x}_i^\top$ and is thus positive semidefinite. By the second order conditions, f_0 is convex! Thus we have turned an extremely non-convex problem into an unconstrained convex minimization problem just by a neat application of monotone functions. We can efficiently solve this problem algorithmically using convex optimization solvers such as gradient descent.

Actually, we can do better than gradient descent for this particular example! If we define

$$X = \begin{bmatrix} \vec{x}_1^\top \\ \vdots \\ \vec{x}_n^\top \end{bmatrix} \in \mathbb{R}^{n \times d}, \quad \vec{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} \in \mathbb{R}^n, \quad \vec{p}(\vec{w}) = \begin{bmatrix} \sigma(\vec{x}_1^\top \vec{w}) \\ \vdots \\ \sigma(\vec{x}_n^\top \vec{w}) \end{bmatrix} \in \mathbb{R}^n \quad (5.196)$$

then the gradient looks like

$$\nabla f_0(\vec{w}) = X^\top (\vec{p}(\vec{w}) - \vec{y}). \quad (5.197)$$

Notice the similarity to the gradient of least squares, $X^\top (X\vec{w} - \vec{y})$. In fact, we can exploit this similarity to obtain a specialized optimization algorithm called *iteratively reweighted least squares*, which can efficiently solve for maximum likelihood estimates within this type of statistical model (i.e., a *generalized linear model*).

5.5.2 Slack Variables

The purpose of slack variables is to turn equality constraints into inequality constraints and vice-versa.

Definition 140 (Slack Variables)

Consider a problem of the form

$$\begin{aligned} \min_{\vec{x} \in \mathbb{R}^n} \quad & f_0(\vec{x}) \\ \text{s.t.} \quad & f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}. \end{aligned} \quad (5.198)$$

Let $\mathcal{S} \subseteq \{1, \dots, m\}$ be any subset. The above problem is equivalent to the following problem:

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{s} \in \mathbb{R}_+^{\mathcal{S}}}} f_0(\vec{x}) \quad (5.199)$$

$$\begin{aligned} \text{s.t. } f_i(\vec{x}) + s_i &= 0, & \forall i \in \mathcal{S} \\ f_i(\vec{x}) &\leq 0, & \forall i \in \{1, \dots, m\} \setminus \mathcal{S} \\ h_j(\vec{x}) &= 0, & \forall j \in \{1, \dots, p\}. \end{aligned}$$

Here the notation $\mathbb{R}_+^{\mathcal{S}} = \{(x_i)_{i \in \mathcal{S}} \mid x_i \geq 0 \forall i \in \mathcal{S}\}$, and \vec{s} is called a *slack variable*.

One can choose to create slack variables s_i for only a subset of the inequality constraints, or all of them. When we work with more advanced optimization algorithms later, sometimes this parameterization is crucial (e.g. for equality-constrained Newton's method).

Example 141. If we have a problem of the form

$$\min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (5.200)$$

$$\text{s.t. } 3x_1^2 + 4x_2^2 \leq 0 \quad (5.201)$$

$$2x_1^2 + 5x_2^2 \leq 0 \quad (5.202)$$

but our solver could only handle equality constraints, then it would be equivalent to solve the problem

$$\min_{\substack{\vec{x} \in \mathbb{R}^2 \\ \vec{s} \in \mathbb{R}_+^2}} f_0(\vec{x}) \quad (5.203)$$

$$\text{s.t. } 3x_1^2 + 4x_2^2 + s_1 = 0 \quad (5.204)$$

$$2x_1^2 + 5x_2^2 + s_2 = 0 \quad (5.205)$$

and, upon solving this problem and obtaining (\vec{x}^*, \vec{s}^*) , the solution to the original problem would be this same \vec{x}^* .

5.5.3 Epigraph Reformulation

The epigraph reformulation is a way to convert between non-linearity in the objective and a constraint.

Definition 142

Consider a problem of the form

$$\min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (5.206)$$

$$\text{s.t. } f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\}$$

$$h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}.$$

Its *epigraph reformulation* is the problem

$$\min_{\substack{t \in \mathbb{R} \\ \vec{x} \in \mathbb{R}^n}} t \quad (5.207)$$

$$\text{s.t. } t \geq f_0(\vec{x})$$

$$f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\}$$

$$h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}.$$

The epigraph objective is always a linear and differentiable function of the decision variables (t, \vec{x}) . However, the constraint can become complicated if $f_0(\vec{x})$ is non-linear. This transformation is especially useful in the case of quadratically-constrained quadratic programs (QCQP).

Example 143 (Elastic-Net Regularization). This example uses the two previously-discussed techniques in tandem to figure out how to handle a regularizer with both smooth and non-smooth components (i.e., ℓ^1 and ℓ^2 norms).

Let $A \in \mathbb{R}^{m \times n}$, and $\vec{y} \in \mathbb{R}^m$. Suppose that we have a problem of the form

$$\min_{\vec{x} \in \mathbb{R}^n} \left\{ \|A\vec{x} - \vec{y}\|_2^2 + \alpha \|\vec{x}\|_2^2 + \beta \|\vec{x}\|_1 \right\}. \quad (5.208)$$

The regularizer $\alpha \|\vec{x}\|_2^2 + \beta \|\vec{x}\|_1$ is called the *elastic net regularizer* and encourages “sparse” and small \vec{x} ; this regularizer has some use in the analysis of high-dimensional and structured data.

Suppose that our solver can only handle differentiable objectives, but is able to handle constraints so long as they are also differentiable. Then we cannot solve the problem out-right using our solver, so we need to reformulate it. We can first start by using a modification of the epigraph reformulation:

$$\min_{\substack{t \in \mathbb{R} \\ \vec{x} \in \mathbb{R}^n}} \|A\vec{x} - \vec{y}\|_2^2 + \alpha \|\vec{x}\|_2^2 + \beta t \quad (5.209)$$

$$\text{s.t. } t \geq \|\vec{x}\|_1. \quad (5.210)$$

Now the constraint is non-differentiable, so we are no longer able to exactly solve this constrained problem. However, the objective is now convex and differentiable. Let us rewrite this problem using the $|x_i|$:

$$\min_{\substack{t \in \mathbb{R} \\ \vec{x} \in \mathbb{R}^n}} \|A\vec{x} - \vec{y}\|_2^2 + \alpha \|\vec{x}\|_2^2 + \beta t \quad (5.211)$$

$$\text{s.t. } t \geq \sum_{i=1}^n |x_i|. \quad (5.212)$$

The main insight that goes into resolving the non-differentiability of this constraint is that $|x_i| = \max\{x_i, -x_i\}$ and in particular $s_i \geq |x_i|$ if and only if $s_i \geq x_i$ and $s_i \geq -x_i$. Thus we add more “slack-type” variables and obtain

$$\min_{\substack{t \in \mathbb{R} \\ \vec{x} \in \mathbb{R}^n \\ \vec{s} \in \mathbb{R}_+^n}} \|A\vec{x} - \vec{y}\|_2^2 + \alpha \|\vec{x}\|_2^2 + \beta t \quad (5.213)$$

$$\text{s.t. } t \geq \sum_{i=1}^n s_i \quad (5.214)$$

$$s_i \geq x_i, \quad \forall i \in \{1, \dots, n\} \quad (5.215)$$

$$s_i \geq -x_i, \quad \forall i \in \{1, \dots, n\}. \quad (5.216)$$

Now the constraints are all differentiable (in fact, affine) and the problem may be solved. This problem is exactly equivalent to the above problem because t is being minimized in the objective, so each s_i is being minimized by way of the first constraint, meaning that $s_i = |x_i|$ and this is equivalent to the original elastic-net regression problem.

Chapter 6

Gradient Descent

Relevant sections of the textbooks:

- [2] Section 12.2.

6.1 Strong Convexity and Smoothness

In this section, we will introduce two properties of functions that will become useful in analyzing optimization algorithms. The first property is a notion of convexity of functions that is stronger than the ones previously introduced.

Definition 144 (μ -Strongly Convex Function)

Let $\mu \geq 0$. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function. We say that f is μ -strongly convex if the domain of f , say Ω , is convex, and, for any $\vec{x}, \vec{y} \in \Omega$, we have


$$f(\vec{y}) \geq f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x}) + \frac{\mu}{2} \|\vec{y} - \vec{x}\|_2^2. \quad (6.1)$$

A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is μ -strongly concave if $-f$ is μ -strongly convex.

Recall from Theorem 123 that the first order condition for (usual) convexity requires the function to be bigger than its linear (first order) Taylor approximation centered at any point. μ -strong convexity imposes a stronger requirement on the function: it needs to be bigger than its linear approximation plus a non-negative quadratic term that has a Hessian matrix μI . This becomes more obvious if we write Equation (6.1) in the equivalent form

$$f(\vec{y}) \geq \underbrace{f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x})}_{\text{first-order Taylor approximation}} + \underbrace{(\vec{y} - \vec{x})^\top \left(\frac{\mu}{2} I \right) (\vec{y} - \vec{x})}_{\text{non-negative quadratic term}}. \quad (6.2)$$

Below, we visualize the μ -strong convexity property and compare it to the first order condition for convexity.



Therefore, μ -strong convexity of the function is a very important feature. It guarantees that the function will always have enough *curvature* (at least as much as its quadratic lower bound) and thus will never become too flat anywhere on its domain.

If the function f is twice-differentiable we can formalize this notion by giving the following equivalent condition for μ -strong convexity.

Theorem 145 (Second Order Condition for μ -Strong Convexity)

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a twice-differentiable function. Then f is μ -strongly convex if and only if for all $\vec{x} \in \Omega$ we have

$$\nabla^2 f(\vec{x}) - \mu I \succeq 0, \quad (6.3)$$

i.e., $\nabla^2 f(\vec{x}) - \mu I$ is PSD for each $\vec{x} \in \Omega$.

An important property of μ -strongly convex functions is that they are strictly convex and thus they have at most one minimizer. In fact, one can show that they have exactly one minimizer.

Theorem 146 (Strongly Convex Functions have Unique Global Minimizers)

Let $\mu > 0$. Let $\Omega \subseteq \mathbb{R}^n$ be a convex set and let $f: \Omega \rightarrow \mathbb{R}$ be a μ -strongly convex function. Then f has *exactly one* global minimizer.

The second property we want to introduce is L -smoothness, which describes quadratic *upper* bounds on the function f .

Definition 147 (L -Smooth Function)

Let $L \geq 0$. Let $\Omega \subseteq \mathbb{R}^n$ be a set. Let $f: \Omega \rightarrow \mathbb{R}$ be a differentiable function. We say that f is L -smooth if, for


any $\vec{x}, \vec{y} \in \Omega$, we have

$$f(\vec{y}) \leq f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x}) + \frac{L}{2} \|\vec{y} - \vec{x}\|_2^2. \quad (6.4)$$

If f is L -smooth, then the function f is *upper* bounded by its first order Taylor approximation plus a non-negative quadratic term:


$$f(\vec{y}) \leq \underbrace{f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x})}_{\text{first-order Taylor approximation}} + \underbrace{(\vec{y} - \vec{x})^\top \left(\frac{L}{2} I \right) (\vec{y} - \vec{x})}_{\text{non-negative quadratic term}}. \quad (6.5)$$

We can visualize the L -smoothness condition similarly to the strongly convex condition, as below:



L -smoothness provides a quadratic upper bound on f . This upper bound ensures that the function doesn't have too much curvature anywhere on its domain (at most as much as its upper bound). We will see later in this chapter that this actually translates into an upper bound on the rate at which the gradient of the function changes.

Finally, we visualize the behavior of the μ -strongly convex and L -smooth bounds together.



6.2 Gradient Descent

We now have all the tools we need to introduce, understand, and analyze gradient descent - one of the most ubiquitous optimization algorithms. The algorithm is used to numerically solve differentiable and unconstrained optimization problems.

More formally, let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function. We attempt to solve the problem:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}). \quad (6.6)$$

The general idea behind the gradient descent algorithm is that it starts with some *initial guess* $\vec{x}_0 \in \mathbb{R}^n$ and produces a sequence of refined guesses $\vec{x}_1, \vec{x}_2, \dots$, called *iterates*. In each iteration $t = 0, 1, 2, \dots$, the algorithm updates its guess according to the following rule:

$$\vec{x}_{t+1} = \vec{x}_t + \eta \vec{v}_t \quad (6.7)$$

for some $\vec{v}_t \in \mathbb{R}^n$ and $\eta \geq 0$.

There are two quantities that we need to specify to complete the definition of the algorithm:

- The vector \vec{v}_t , or the *search direction*, which specifies a good direction to move.
- The scalar η , or the *step size*, which specifies how far we move in the direction of \vec{v}_t .

For the gradient descent algorithm, we assume that at every point $\vec{x} \in \mathbb{R}^n$ we can get two pieces of information about the function we are optimizing: the value of the function $f(\vec{x}) \in \mathbb{R}$ as well as its gradient $\nabla f(\vec{x}) \in \mathbb{R}^n$. Next, we will use this available information to come up with a good search direction \vec{v}_t . The choice of the step size η is a more difficult task. There is no universal choice of η that is good for all problems and a good choice of η is problem-specific. We will discuss the choice of the step size later in the section and show the important role it plays in the algorithm.

6.2.1 Search Direction

To choose \vec{v}_t , we want to ensure that $f(\vec{x}_t + \eta\vec{v}_t) \leq f(\vec{x}_t)$ for some small η . This motivates taking \vec{v}_t as the *direction of steepest descent*, i.e., the direction in which f decays the fastest. The degree (i.e., the steepness) of the rate of change can be characterized by the directional derivative $Df(\vec{x})[\vec{v}] = \vec{v}^\top [\nabla f(\vec{x})]$.

Theorem 148 (Negative Gradient is Direction of Steepest Descent)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function, and let $\vec{x} \in \mathbb{R}^n$. Then

$$-\frac{\nabla f(\vec{x})}{\|\nabla f(\vec{x})\|_2} \in \operatorname{argmin}_{\substack{\vec{v} \in \mathbb{R}^n \\ \|\vec{v}\|_2=1}} Df(\vec{x})[\vec{v}]. \quad (6.8)$$

Proof. Left as exercise. □

We also want to use the norm of the gradient in our update. For example, if $f(x) = x^2$ then $f'(x) = 2x$, which has large norm when x is far from the optimal point $x^* = 0$. In this way, if the gradient is large then we are usually far away from an optimum, and we want our update $\eta\vec{v}_t$ to be large. This motivates choosing $\vec{v}_t = -\nabla f(\vec{x}_t)$.

6.2.2 Defining Gradient Descent

This choice of \vec{v}_t gives the famous gradient descent iteration scheme:

$$\vec{x}_{t+1} = \vec{x}_t - \eta \nabla f(\vec{x}_t), \quad \forall t = 0, 1, 2, \dots \quad (6.9)$$

We can formalize this in an algorithm which terminates after T iterations for a user-set T .

Algorithm 3 Gradient Descent.

```

1: function GRADIENTDESCENT( $f, \vec{x}_0, \eta, T$ )
2:   for  $t = 0, 1, \dots, T - 1$  do
3:      $\vec{x}_{t+1} \leftarrow \vec{x}_t - \eta \nabla f(\vec{x}_t)$ 
4:   end for
5:   return  $\vec{x}_T$ 
6: end function

```

It is important to note that with the choice $\vec{v}_t = -\nabla f(\vec{x}_t)$, *descent is not guaranteed*. Namely, for a fixed $\eta > 0$, it is *not* true in general that

$$f(\vec{x}_{t+1}) = f(\vec{x}_t - \eta \nabla f(\vec{x}_t)) \leq f(\vec{x}_t). \quad (6.10)$$

Rather, from the proof of the above theorem, we only have that, for any given t there *exists* $\eta_t > 0$ such that

$$f(\vec{x}_t - \eta_t \nabla f(\vec{x}_t)) \leq f(\vec{x}_t). \quad (6.11)$$

This η_t , in general, is very small, and depends heavily on t and the local geometry of f around \vec{x}_t . None of these details are known by any realistic implementation of the gradient descent algorithm. In what follows, we will study gradient descent with a constant step size; this setting is most common in practice.

6.2.3 Convergence Analysis of Gradient Descent

When applying GD algorithm to a problem we need to make sure that we're making progress and that we eventually converge to the optimal solution. This question is very related to the choice of the step size. For the remainder of the section we will show that, for certain classes of functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and certain chosen step sizes $\eta > 0$, the gradient descent algorithm is guaranteed to make progress in every iteration and converge to the optimal solution. We will first explore these questions through a familiar example.

Example 149 (Gradient Descent for Least Squares). In this example we explore the convergence properties of the gradient descent algorithm by applying it to the least squares problem. Let $A \in \mathbb{R}^{m \times n}$ have full column rank, and $\vec{y} \in \mathbb{R}^m$.

$$\min_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{y}\|_2^2. \quad (6.12)$$

Recall that, in this setting, the least squares problem has the unique closed form solution

$$\vec{x}^* = (A^\top A)^{-1} A^\top \vec{y}. \quad (6.13)$$

We will use our knowledge of the true solution to analyze the convergence of the gradient descent algorithm. To apply gradient descent to this problem, let us first compute the gradient, which we see as

$$\nabla f(\vec{x}) = 2A^\top (A\vec{x} - \vec{y}). \quad (6.14)$$

Now let $\vec{x}_0 \in \mathbb{R}^n$ be the initial guess. For t a non-negative integer, we can write the gradient descent step:

$$\vec{x}_{t+1} = \vec{x}_t - \eta \nabla f(\vec{x}_t) \quad (6.15)$$

$$= \vec{x}_t - 2\eta A^\top (A\vec{x}_t - \vec{y}) \quad (6.16)$$

$$= (I - 2\eta A^\top A)\vec{x}_t + 2\eta A^\top \vec{y}. \quad (6.17)$$

We aim to set η which achieves the following two desired properties for the gradient descent iterates \vec{x}_t :

- We make progress, i.e., in every iteration we get closer to the optimal solution:

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2 \leq \|\vec{x}_t - \vec{x}^*\|_2. \quad (6.18)$$

- We eventually converge to the optimal solution:

$$\lim_{t \rightarrow \infty} \vec{x}_t = \vec{x}^* \iff \lim_{t \rightarrow \infty} \|\vec{x}_t - \vec{x}^*\|_2 = 0. \quad (6.19)$$

To study this, let us write out the relationship between $\vec{x}_{t+1} - \vec{x}^*$ and $\vec{x}_t - \vec{x}^*$. We do this by subtracting \vec{x}^* from both sides of Equation (6.17) and do some algebraic manipulations.

$$\vec{x}_{t+1} - \vec{x}^* = (I - 2\eta A^\top A)\vec{x}_t + 2\eta A^\top \vec{y} - \vec{x}^* \quad (6.20)$$

$$= (I - 2\eta A^\top A)\vec{x}_t + 2\eta A^\top \vec{y} - (A^\top A)^{-1} A^\top \vec{y} \quad (6.21)$$

$$= (I - 2\eta A^\top A)\vec{x}_t + 2\eta (A^\top A)(A^\top A)^{-1} A^\top \vec{y} - (A^\top A)^{-1} A^\top \vec{y} \quad (6.22)$$

$$= (I - 2\eta A^\top A)\vec{x}_t + 2\eta A^\top A\vec{x}^* - \vec{x}^* \quad (6.23)$$

$$= (I - 2\eta A^\top A)\vec{x}_t - (I - 2\eta A^\top A)\vec{x}^* \quad (6.24)$$

$$= (I - 2\eta A^\top A)(\vec{x}_t - \vec{x}^*). \quad (6.25)$$

Here we used the trick of introducing $I = (A^\top A)(A^\top A)^{-1}$; this is because we wanted to group terms with $A^\top A$, and we also wanted to introduce instances of $\vec{x}^* = (A^\top A)^{-1}A^\top \vec{y}$. Thus, while this was a trick, it was a motivated trick. In the end we get the following relationship:

$$\vec{x}_{t+1} - \vec{x}^* = (I - 2\eta A^\top A)(\vec{x}_t - \vec{x}^*). \quad (6.26)$$

If we take the norm of both sides, we have

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2 = \|(I - 2\eta A^\top A)(\vec{x}_t - \vec{x}^*)\|_2 \quad (6.27)$$

$$\leq \|I - 2\eta A^\top A\|_2 \|\vec{x}_t - \vec{x}^*\|_2 \quad (6.28)$$

$$= \sigma_{\max}\{I - 2\eta A^\top A\} \|\vec{x}_t - \vec{x}^*\|_2. \quad (6.29)$$

Thus, if $\sigma_{\max}\{I - 2\eta A^\top A\} < 1$, then we have

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2 \leq \sigma_{\max}\{I - 2\eta A^\top A\} \|\vec{x}_t - \vec{x}^*\|_2 < \|\vec{x}_t - \vec{x}^*\|_2. \quad (6.30)$$

Thus we are guaranteed to make progress at each step. For the convergence guarantee, we recursively apply Equation (6.29) to get

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2 \leq \sigma_{\max}\{I - 2\eta A^\top A\} \|\vec{x}_t - \vec{x}^*\|_2 \quad (6.31)$$

$$\leq \sigma_{\max}\{I - 2\eta A^\top A\}^2 \|\vec{x}_{t-1} - \vec{x}^*\|_2 \quad (6.32)$$

$$\leq \dots \quad (6.33)$$

$$\leq \sigma_{\max}\{I - 2\eta A^\top A\}^{t+1} \|\vec{x}_0 - \vec{x}^*\|_2. \quad (6.34)$$

Thus taking the limit we get

$$\lim_{t \rightarrow \infty} \|\vec{x}_t - \vec{x}^*\|_2 \leq \lim_{t \rightarrow \infty} \sigma_{\max}\{I - 2\eta A^\top A\}^t \|\vec{x}_0 - \vec{x}^*\|_2 \quad (6.35)$$

$$= \|\vec{x}_0 - \vec{x}^*\|_2 \cdot \underbrace{\lim_{t \rightarrow \infty} \sigma_{\max}\{I - 2\eta A^\top A\}^t}_{=0} \quad (6.36)$$

$$= 0. \quad (6.37)$$

Thus our convergence guarantee holds as well, so long as $\sigma_1\{I - 2\eta A^\top A\} < 1$.

Let us now generalize this analysis to a larger class of functions that includes least squares, as well as more general functions. We will focus our attention on the class of L -smooth and μ -strongly convex functions. Similar to what we did for least squares, we will use the optimal solution \vec{x}^* in our analysis. For a general function f , the quantity \vec{x}^* is almost always not known. But in the case of L -smooth and μ -strongly convex functions, a unique global minimizer exists. We just need the fact that it exists to show that for some small enough choice of the step size η , the gradient descent algorithm converges to this optimal solution. Before we formally state and prove this, we want to introduce one property of L -smooth functions that will become useful in our following proof.

Lemma 150 (Gradient Bound for L -Smooth Functions)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be an L -smooth function. For all $\vec{x} \in \mathbb{R}^n$, we have

$$\|\nabla f(\vec{x})\|_2^2 \leq 2L \left(f(\vec{x}) - \min_{\vec{x}' \in \mathbb{R}^n} f(\vec{x}') \right). \quad (6.38)$$

This lemma says that the magnitude of the gradient gets smaller as we get closer to the optimal solution.

Proof. Recall the definition of L -smooth functions:

$$f(\vec{y}) \leq f(\vec{x}) + [\nabla f(\vec{x})]^\top (\vec{y} - \vec{x}) + \frac{L}{2} \|\vec{y} - \vec{x}\|_2^2. \quad (6.39)$$

This is true for all points $\vec{x}, \vec{y} \in \mathbb{R}^n$, so let us fix \vec{x} and set $\vec{y} = \vec{x} - \frac{\nabla f(\vec{x})}{L}$. We get

$$f\left(\vec{x} - \frac{\nabla f(\vec{x})}{L}\right) \leq f(\vec{x}) + [\nabla f(\vec{x})]^\top \left(-\frac{\nabla f(\vec{x})}{L}\right) + \frac{L}{2} \left\|-\frac{\nabla f(\vec{x})}{L}\right\|_2^2 \quad (6.40)$$

$$= f(\vec{x}) - \frac{1}{L} \|\nabla f(\vec{x})\|_2^2 + \frac{1}{2L} \|\nabla f(\vec{x})\|_2^2 \quad (6.41)$$

$$= f(\vec{x}) - \frac{1}{2L} \|\nabla f(\vec{x})\|_2^2. \quad (6.42)$$

Now we can use the fact that $\min_{\vec{x}'} f(\vec{x}') \leq f(\vec{z})$ for all $\vec{z} \in \mathbb{R}^n$ to get

$$\min_{\vec{x}' \in \mathbb{R}^n} f(\vec{x}') \leq f\left(\vec{x} - \frac{\nabla f(\vec{x})}{L}\right) \quad (6.43)$$

$$\leq f(\vec{x}) - \frac{1}{2L} \|\nabla f(\vec{x})\|_2^2. \quad (6.44)$$

Rearranging, we have

$$\|\nabla f(\vec{x})\|_2^2 \leq 2L \left(f(\vec{x}) - \min_{\vec{x}' \in \mathbb{R}^n} f(\vec{x}') \right) \quad (6.45)$$

as desired. \square

Now we have all the needed tools to prove the following property of gradient descent: for any μ -strongly convex and L -smooth function, the gradient descent algorithm converges exponentially fast to the true solution \vec{x}^* .

Theorem 151 (Convergence of Gradient Descent for Smooth Strongly Convex Functions)

Let $\mu, L > 0$. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be an L -smooth, μ -strongly convex function. Consider the following optimization problem:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}) \quad (6.46)$$

which has optimal solution \vec{x}^* . Then the constant step size $\eta = \frac{1}{L}$ is such that, if applying the gradient descent algorithm generates the sequence of points

$$\vec{x}_{t+1} = \vec{x}_t - \eta \nabla f(\vec{x}_t), \quad (6.47)$$

then for all initializations \vec{x}_0 and non-negative integers t ,

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2^2 \leq \left(1 - \frac{\mu}{L}\right) \|\vec{x}_t - \vec{x}^*\|_2^2. \quad (6.48)$$

Proof. We want to upper bound $\|\vec{x}_{t+1} - \vec{x}^*\|_2^2$, so let us write it out as

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2^2 = \|\vec{x}_t - \eta \nabla f(\vec{x}_t) - \vec{x}^*\|_2^2 \quad (6.49)$$

$$= \|[\vec{x}_t - \vec{x}^*] - \eta \nabla f(\vec{x}_t)\|_2^2 \quad (6.50)$$

$$= \|\vec{x}_t - \vec{x}^*\|_2^2 - 2\eta [\nabla f(\vec{x}_t)]^\top (\vec{x}_t - \vec{x}^*) + \eta^2 \|\nabla f(\vec{x}_t)\|_2^2 \quad (6.51)$$

$$= \|\vec{x}_t - \vec{x}^*\|_2^2 + 2\eta[\nabla f(\vec{x}_t)]^\top(\vec{x}^* - \vec{x}_t) + \eta^2 \|\nabla f(\vec{x}_t)\|_2^2. \quad (6.52)$$

Because f is L -smooth, the lemma gives

$$\|\nabla f(\vec{x}_t)\|_2^2 \leq 2L(f(\vec{x}_t) - f(\vec{x}^*)). \quad (6.53)$$

Because f is μ -strongly convex, we can write

$$f(\vec{x}^*) \geq f(\vec{x}_t) + [\nabla f(\vec{x}_t)]^\top(\vec{x}^* - \vec{x}_t) + \frac{\mu}{2} \|\vec{x}^* - \vec{x}_t\|_2^2 \quad (6.54)$$

$$\implies [\nabla f(\vec{x}_t)]^\top(\vec{x}_t - \vec{x}^*) \geq f(\vec{x}_t) - f(\vec{x}^*) + \frac{\mu}{2} \|\vec{x}_t - \vec{x}^*\|_2^2 \quad (6.55)$$

$$\implies [\nabla f(\vec{x}_t)]^\top(\vec{x}^* - \vec{x}_t) \leq f(\vec{x}^*) - f(\vec{x}_t) - \frac{\mu}{2} \|\vec{x}_t - \vec{x}^*\|_2^2. \quad (6.56)$$

Plugging these two estimates in, we get

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2^2 = \|\vec{x}_t - \vec{x}^*\|_2^2 + 2\eta[\nabla f(\vec{x}_t)]^\top(\vec{x}^* - \vec{x}_t) + \eta^2 \|\nabla f(\vec{x}_t)\|_2^2 \quad (6.57)$$

$$\leq \|\vec{x}_t - \vec{x}^*\|_2^2 + 2\eta \left[f(\vec{x}^*) - f(\vec{x}_t) - \frac{\mu}{2} \|\vec{x}_t - \vec{x}^*\|_2^2 \right] + \eta^2 [2L(f(\vec{x}_t) - f(\vec{x}^*))] \quad (6.58)$$

$$= (1 - \eta\mu) \|\vec{x}_t - \vec{x}^*\|_2^2 + 2\eta(\eta L - 1)(f(\vec{x}_t) - f(\vec{x}^*)). \quad (6.59)$$

So as to make the second term 0, we choose $\eta = \frac{1}{L}$. This gives

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2^2 \leq \left(1 - \frac{\mu}{L}\right) \|\vec{x}_t - \vec{x}^*\|_2^2. \quad (6.60)$$

□

Corollary 152. Consider the setting of Theorem 151. We have that $0 \leq 1 - \frac{\mu}{L} < 1$, and consequently:

(a) (Descent at every step.) $\|\vec{x}_{t+1} - \vec{x}^*\|_2 \leq \|\vec{x}_t - \vec{x}^*\|_2$ for all non-negative integers t and initializations \vec{x}_0 .

(b) (Convergence.) $\lim_{t \rightarrow \infty} \vec{x}_t = \vec{x}^*$.

Proof. We need to show that $1 - \frac{\mu}{L} \in [0, 1)$, or in other words that $0 < \frac{\mu}{L} \leq 1$. By assumption $\mu > 0$ so $\frac{\mu}{L} > 0$. The upper bound is given by the fact that f is μ -strongly convex and L -smooth, so

$$f(\vec{x}) + [\nabla f(\vec{x})]^\top(\vec{y} - \vec{x}) + \frac{\mu}{2} \|\vec{y} - \vec{x}\|_2^2 \leq f(\vec{y}) \leq f(\vec{x}) + [\nabla f(\vec{x})]^\top(\vec{y} - \vec{x}) + \frac{L}{2} \|\vec{y} - \vec{x}\|_2^2 \quad (6.61)$$

which implies that

$$f(\vec{x}) + [\nabla f(\vec{x})]^\top(\vec{y} - \vec{x}) + \frac{\mu}{2} \|\vec{y} - \vec{x}\|_2^2 \leq f(\vec{x}) + [\nabla f(\vec{x})]^\top(\vec{y} - \vec{x}) + \frac{L}{2} \|\vec{y} - \vec{x}\|_2^2 \quad (6.62)$$

which implies that

$$\frac{\mu}{2} \|\vec{y} - \vec{x}\|_2^2 \leq \frac{L}{2} \|\vec{y} - \vec{x}\|_2^2 \quad (6.63)$$

which finally implies that $\mu \leq L$, i.e., $0 < \frac{\mu}{L} \leq 1$. This immediately yields the first claim since $1 - \frac{\mu}{L} \in [0, 1)$ so $\sqrt{1 - \frac{\mu}{L}} \in [0, 1)$, therefore

$$\|\vec{x}_{t+1} - \vec{x}^*\|_2 \leq \sqrt{1 - \frac{\mu}{L}} \|\vec{x}_t - \vec{x}^*\|_2 < \|\vec{x}_t - \vec{x}^*\|_2. \quad (6.64)$$

The second claim follows since

$$\|\vec{x}_t - \vec{x}^*\|_2 \leq \left(1 - \frac{\mu}{L}\right) \|\vec{x}_{t-1} - \vec{x}^*\|_2 \quad (6.65)$$

$$\leq \left(1 - \frac{\mu}{L}\right)^2 \|\vec{x}_{t-2} - \vec{x}^*\|_2^2 \quad (6.66)$$

$$\leq \dots \quad (6.67)$$

$$\leq \left(1 - \frac{\mu}{L}\right)^t \|\vec{x}_0 - \vec{x}^*\|_2^2 \quad (6.68)$$

and so

$$\lim_{t \rightarrow \infty} \|\vec{x}_t - \vec{x}^*\|_2^2 = \lim_{t \rightarrow \infty} \left(1 - \frac{\mu}{L}\right)^t \|\vec{x}_0 - \vec{x}^*\|_2^2 \quad (6.69)$$

$$= \|\vec{x}_0 - \vec{x}^*\|_2^2 \cdot \underbrace{\lim_{t \rightarrow \infty} \left(1 - \frac{\mu}{L}\right)^t}_{=0} \quad (6.70)$$

$$= 0. \quad (6.71)$$

□

In this case the quantity $c = \sqrt{1 - \frac{\mu}{L}}$ is important. If we wanted to run gradient descent for T iterations to within accuracy ϵ , we would have

$$\|\vec{x}_T - \vec{x}^*\|_2 \leq c^T \|\vec{x}_0 - \vec{x}^*\|_2. \quad (6.72)$$

If we set $D \doteq \|\vec{x}_0 - \vec{x}^*\|$, then we can ensure that $\|\vec{x}_T - \vec{x}^*\|_2 \leq \epsilon$ by bounding the right-hand side $c^T \|\vec{x}_0 - \vec{x}^*\|_2 = c^T D$ by ϵ , getting

$$c^T D \leq \epsilon \quad (6.73)$$

$$\implies c^T \leq \frac{\epsilon}{D} \quad (6.74)$$

$$\implies T \log(c) \leq \log(\epsilon) - \log(D) \quad (6.75)$$

$$\implies T \geq \frac{\log(\epsilon) - \log(D)}{\log(c)} \quad (6.76)$$

$$= \frac{\log(1/\epsilon) + \log(D)}{\log(1/c)}. \quad (6.77)$$

Thus, the lower the value of c is, the faster we get convergence towards a given accuracy.

Finally, it is important to point out that this is not the only class of functions for which the gradient descent algorithm converges. For example for the class of functions that are L -smooth and convex (i.e., not μ -strictly convex), the gradient descent algorithm does converge; in particular it converges to within accuracy ϵ after $T \geq O(1/\epsilon)$ iterations. This is vastly slower than the convergence achieved for μ -strongly convex functions.

6.3 Variations: Stochastic Gradient Descent

In this section we will cover the stochastic gradient descent (SGD) algorithm. This variant is motivated by optimization problems in which computing the gradient can be computationally expensive. The idea behind SGD is to use random sampling to find an estimate of the gradient that is easy to compute. Let us consider the class of unconstrained differentiable optimization problems that take the form

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}) = \min_{\vec{x} \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m f_i(\vec{x}). \quad (6.78)$$

The gradient of the objective function is given by

$$\nabla f(\vec{x}) = \frac{1}{m} \sum_{i=1}^m \nabla f_i(\vec{x}), \quad (6.79)$$

and the gradient descent step is given by

$$\vec{x}_{t+1} = \vec{x}_t - \eta \cdot \frac{1}{m} \sum_{i=1}^m \nabla f_i(\vec{x}), \quad (6.80)$$

This class of functions is very common in applications that involve learning from data, one example being the least squares problem. Recall that we can write the least squares problem as

$$\min_{\vec{x} \in \mathbb{R}^n} \underbrace{\frac{1}{m} \|A\vec{x} - \vec{y}\|_2^2}_{=f(\vec{x})} = \min_{\vec{x} \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m \underbrace{(\vec{a}_i^\top \vec{x} - b_i)^2}_{=f_i(\vec{x})} \quad (6.81)$$

Note that we multiplied the objective function with the constant $\frac{1}{m}$, which does not change the solutions to the optimization problem. We will see shortly that this constant is important in the SGD setting.

If the number of functions m is large and if each gradient $\nabla f_i(\vec{x})$ is complex to evaluate, computing the full gradient in Equation (6.79) can become prohibitively expensive. To overcome this, SGD proposes to take a random sample from the set of functions $\{f_1, f_2, \dots, f_m\}$, evaluate the gradient at that sample only, and take the step

$$\vec{x}_{t+1} = \vec{x}_t - \eta \nabla f_i(\vec{x}). \quad (6.82)$$

If this index i is drawn uniformly at random, we can see that the expected value of the estimated gradient is equal to the full gradient.

$$\mathbb{E}[\nabla f_i(\vec{x})] = \sum_{i=1}^m \frac{1}{m} \nabla f_i(\vec{x}) \quad (6.83)$$

$$= \nabla f(\vec{x}). \quad (6.84)$$

Note that the direction $-\nabla f_i(\vec{x})$ is not guaranteed to be the direction of steepest descent of the function $f(\vec{x})$. In fact, it might not be a descent direction at all, and the value of the function f might even *increase* by taking a step in this direction. Further, it is not guaranteed to satisfy the first order optimality conditions; that is, when $\nabla f(\vec{x}) = \vec{0}$, the gradient of a single function $\nabla f_i(\vec{x})$ is generally not zero. This already tells us that the notion of convergence we studied (which is called “last-iterate convergence”) is practically impossible to achieve — and in particular, the type of convergence proof that we studied for gradient descent will not work here. Instead the type of convergence guarantees we can hope for in this setting is for the *averaged* point across all iterations will converge to the optimal solution, i.e.,

$$\lim_{T \rightarrow \infty} f\left(\frac{1}{T} \sum_{t=1}^T \vec{x}_t\right) = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}). \quad (6.85)$$

Further, to prove convergence of SGD we need to use a variable step size η_t . Using a fixed step size (like we did for gradient descent) doesn’t guarantee convergence. The intuition behind this is that the gradient directions $\nabla f_i(\vec{x})$ for different i might be competing with each other and want to pull the solution in different directions. This will cause the sequence \vec{x}_t to bounce back and forth. Using a variable step size η_t such that $\lim_{t \rightarrow \infty} \eta_t = 0$ makes it possible to converge to a single point. It is not trivial to show that the point that the algorithm will converge to is the optimal solution, but this can be proven under some assumptions on the objective function. For formal proofs, a good reference is [6].

Example 153 (Finding the Centroid Using SGD). Fix points $\vec{p}_1, \dots, \vec{p}_m \in \mathbb{R}^n$. Let us consider the problem of finding their *centroid*, which we formulate as

$$\min_{\vec{x} \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m \underbrace{\frac{1}{2} \|\vec{x} - \vec{p}_i\|_2^2}_{=f_i(\vec{x})}. \quad (6.86)$$

The optimal solution for this problem is just the average of the points

$$\vec{x}^* = \frac{1}{m} \sum_{i=1}^m \vec{p}_i. \quad (6.87)$$

Let us apply SGD on this problem with the time varying step size $\eta_t = \frac{1}{t}$ and initial guess $\vec{x}_0 = \vec{0}$. We first compute the gradients

$$\nabla f_i(\vec{x}) = \vec{x} - \vec{p}_i. \quad (6.88)$$

To simplify the notation we will apply SGD by selecting the indices in order.

$$\text{Iteration 1 : } \vec{x}_1 = \vec{x}_0 - \frac{1}{1}(\vec{x}_0 - \vec{p}_1) = \vec{p}_1 \quad (6.89)$$

$$\text{Iteration 2 : } \vec{x}_2 = \vec{x}_1 - \frac{1}{2}(\vec{x}_1 - \vec{p}_2) = \frac{\vec{p}_1 + \vec{p}_2}{2} \quad (6.90)$$

$$\text{Iteration 3 : } \vec{x}_3 = \vec{x}_2 - \frac{1}{3}(\vec{x}_2 - \vec{p}_3) = \frac{\vec{p}_1 + \vec{p}_2 + \vec{p}_3}{3} \quad (6.91)$$

$$\vdots \quad (6.92)$$

$$\text{Iteration } m : \vec{x}_m = \vec{x}_{m-1} - \frac{1}{m}(\vec{x}_{m-1} - \vec{p}_m) = \frac{\sum_{i=1}^m \vec{p}_i}{m}. \quad (6.93)$$

So the SGD algorithm which takes a step along the gradient of a single $\nabla f_i(\vec{x})$ in every iteration converges to the true solution in m iterations. It is important to note that this is a toy example that is used to illustrate the application of SGD. As we discussed earlier, this type of convergence behavior is not common or guaranteed when applying SGD to more complicated real world problems.

Finally, we note that this set up of the SGD algorithm allows us to think of the optimization scheme as an online algorithm. Instead of randomly selecting an index to compute the gradient, assume that the data points (i.e., in the above example the \vec{p}_i) that define the functions $f_i(\vec{x})$ arrive one at a time. Following the same idea of SGD, we can perform an optimization step as each new data point becomes available. In this setting it is important to normalize by the number of data points m in the objective function; otherwise the objective function will keep growing as we get more data, irregardless of whether our optimization algorithm finds a good solution.

6.4 Variations: Gradient Descent for Constrained Optimization

So far we have seen how gradient descent can be applied to different problems all of which are unconstrained. There are variants of the gradient descent algorithm that can be used to solve problems with *simple* constraints. Let $\Omega \subseteq \mathbb{R}^n$ be a *convex* set, let $f: \Omega \rightarrow \mathbb{R}$ be a differentiable function, and consider the problem

$$p^* = \min_{\vec{x} \in \Omega} f(\vec{x}). \quad (6.94)$$

For these types of problems we want to limit our search to points inside the set Ω . If we start with an initial guess $\vec{x}_0 \in \Omega$ and apply the (unconstrained) gradient descent algorithm

$$\vec{x}_{t+1} = \vec{x}_t - \eta \nabla f(\vec{x}_t), \quad (6.95)$$

the point \vec{x}_{t+1} might end up outside of the feasible set Ω , even when \vec{x}_t is in Ω . In this section we will consider two variants of the gradient descent algorithm that propose techniques to deal with constraints.

6.4.1 Projected Gradient Descent

This first variant is projected gradient descent, proposes the idea of *projecting* the point back into the feasible set.

Definition 154 (Projection onto a Convex Set)

Let Ω be a closed^a convex set. Let $\vec{y} \in \mathbb{R}^n$ be any vector. We define the projection of the vector \vec{y} onto the set Ω as

$$\text{proj}_{\Omega}(\vec{y}) = \underset{\vec{x} \in \Omega}{\text{argmin}} \|\vec{x} - \vec{y}\|_2^2. \quad (6.96)$$

^aA closed set is one that contains its boundary points.

One can show that, since Ω is closed and convex, the projection is unique. In words, $\text{proj}_{\Omega}(\vec{y})$ is the closest point in Ω to \vec{y} . From this definition one sees that if $\vec{y} \in \Omega$ then $\text{proj}_{\Omega}(\vec{y}) = \vec{y}$. Using this definition we can write the step of the projected gradient descent algorithm as

$$\vec{x}_{t+1} = \text{proj}_{\Omega}(\vec{x}_t - \eta \nabla f(\vec{x}_t)). \quad (6.97)$$

Algorithm 4 Projected Gradient Descent

```

1: function PROJECTEDGRADIENTDESCENT( $f, \vec{x}_0, \Omega, \eta, T$ )
2:   for  $t = 0, 1, \dots, T - 1$  do
3:      $\vec{x}_{t+1} \leftarrow \text{proj}_{\Omega}(\vec{x}_t - \eta \nabla f(\vec{x}_t))$ 
4:   end for
5:   return  $\vec{x}_T$ 
6: end function

```

Note that the projection operator itself solves a convex optimization problem. It is only meaningful to consider the projected gradient descent algorithm if this projection problem is simple enough to be solved in every iteration. We have seen examples of projections onto simple sets. For example, the least squares problem computes the projection of a vector onto the subspace $\mathcal{R}(A)$, and we know how to efficiently solve this projection problem. However, that's not always the case, and the projection problem might be nearly as difficult to solve as our original optimization problem.

6.4.2 Conditional Gradient Descent

We introduce another variant of the gradient descent algorithm for constrained problems called conditional gradient descent (also known as Frank-Wolfe method). Recall that in our development of the choice of the “best” search direction for the (regular) gradient descent algorithm, we considered the quantity

$$D_{\vec{v}_t} f(\vec{x}_t) = \lim_{\eta \rightarrow 0} \frac{f(\vec{x}_t + \eta \vec{v}_t) - f(\vec{x}_t)}{\eta} = [\nabla f(\vec{x}_t)]^\top \vec{v}_t. \quad (6.98)$$

This quantity, known as the directional derivative, is the (instantaneous) rate of change of the function f at point \vec{x}_t along direction \vec{v}_t . We also saw that the choice $\vec{v}_t = -\nabla f(\vec{x}_t)$ is the direction that minimizes this rate of change. This is a reasonable choice for the unconstrained case since we can freely move in any direction. The idea behind the

conditional gradient descent algorithm is to limit our search direction to the feasible set Ω and find a vector inside the set that minimizes $\nabla [f(\vec{x}_t)]^\top \vec{v}_t$. Thus, given $\vec{x}_t \in \Omega$, we can define the search direction of the conditional gradient descent as

$$\vec{v}_t = \operatorname{argmin}_{\vec{v} \in \Omega} [\nabla f(\vec{x}_t)]^\top \vec{v}. \quad (6.99)$$

Once we find this direction \vec{v}_t , we need to use it in a way that ensures that we don't leave the set. Taking a step along this direction $\vec{x}_{t+1} = \vec{x}_t + \eta \vec{v}_t$ doesn't guarantee this. But we can do something different and take the convex combination between the current point and the search direction. That is, for $\delta_t \in [0, 1]$ we can update

$$\vec{x}_{t+1} = (1 - \delta_t)\vec{x}_t + \delta_t \vec{v}_t \quad (6.100)$$

By convexity of the set, this guarantees that if we start with a feasible initial guess $\vec{x}_0 \in \Omega$, we get a sequence of points $\vec{x}_t \in \Omega$ for all $t \geq 0$. While this property holds for any choice of $\delta_t \in [0, 1]$, to prove convergence of the algorithm we require $\lim_{t \rightarrow \infty} \delta_t = 0$; a conventional choice is $\delta_t = \frac{1}{t}$.

Algorithm 5 Conditional Gradient Descent (Frank-Wolfe)

```

1: function CONDITIONALGRADIENTDESCENT( $f, \vec{x}_0, \Omega, \delta, T$ )
2:   for  $t = 0, 1, \dots, T - 1$  do
3:      $\vec{v}_t \leftarrow \operatorname{argmin}_{\vec{v} \in \Omega} [\nabla f(\vec{x}_t)]^\top \vec{v}$ 
4:      $\vec{x}_{t+1} \leftarrow (1 - \delta_t)\vec{x}_t + \delta_t \vec{v}_t$ 
5:   end for
6:   return  $\vec{x}_T$ 
7: end function

```

As a last note, we point out that the problem of finding a search direction \vec{v}_t is a constrained optimization problem within Ω .

Chapter 7

Duality

Relevant sections of the textbooks:

- [1] Chapter 5.
- [2] Section 8.5.

7.1 Lagrangian

This chapter develops the theory of duality for optimization problems. This is a technique that can help us solve or bound the optimal values for constrained optimization problems by solving the related dual problem. The dual problem might not always give us a direct solution to our original (“primal”) problem, but it will always give us a bound. For this, we first define the Lagrangian for a problem.

Let us start with our generic constrained optimization problem, which we will call problem \mathcal{P} (which stands for *primal*):

$$\begin{aligned} \text{problem } \mathcal{P}: \quad p^* &= \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) & (7.1) \\ \text{s.t.} \quad f_i(\vec{x}) &\leq 0, & \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) &= 0, & \forall j \in \{1, \dots, p\}. \end{aligned}$$

Let us denote its feasible set by

$$\Omega \doteq \left\{ \vec{x} \in \mathbb{R}^n \mid \begin{array}{l} f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \\ h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\} \end{array} \right\} \quad \text{so that} \quad p^* = \min_{\vec{x} \in \Omega} f_0(\vec{x}). \quad (7.2)$$

We know how to algorithmically approach unconstrained problems, say for instance using gradient descent, which we discussed in the previous chapter. Thus, our question is whether there is a way to incorporate the constraints of the problem \mathcal{P} into the objective function itself. Trying to understand if this is possible is one motivation for the Lagrangian.

To this end, we construct the *indicator function* $\mathbb{1}[\cdot]$, which for a condition $C(\vec{x})$ defined as

$$\mathbb{1}[C(\vec{x})] \doteq \begin{cases} 0, & \text{if } C(\vec{x}) \text{ is true,} \\ +\infty, & \text{if } C(\vec{x}) \text{ is false.} \end{cases} \quad (7.3)$$

Then the problem \mathcal{P} in (7.1) is equivalent to the following *unconstrained* problem:

$$p^* = \min_{\vec{x} \in \Omega} f_0(\vec{x}) \quad (7.4)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \{f_0(\vec{x}) + \mathbb{1}[\vec{x} \in \Omega]\} \quad (7.5)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ f_0(\vec{x}) + \sum_{i=1}^m \mathbb{1}[f_i(\vec{x}) \leq 0] + \sum_{j=1}^p \mathbb{1}[h_j(\vec{x}) = 0] \right\}. \quad (7.6)$$

To explain the chain of equalities leading to (7.6), say we consider the function

$$F_0(\vec{x}) \doteq f_0(\vec{x}) + \sum_{i=1}^m \mathbb{1}[f_i(\vec{x}) \leq 0] + \sum_{j=1}^p \mathbb{1}[h_j(\vec{x}) = 0] \quad (7.7)$$

evaluated at an $\vec{x} \notin \Omega$, i.e., there is some i such that $f_i(\vec{x}) > 0$ or some j such that $h_j(\vec{x}) \neq 0$. Then $F_0(\vec{x}) = \infty$. Therefore no solution to the minimization in (7.6) will ever be outside of Ω , i.e., all solutions to (7.6) will be contained in Ω . And for $\vec{x} \in \Omega$, we have $F_0(\vec{x}) = f_0(\vec{x})$, so that $\min_{\vec{x} \in \Omega} f_0(\vec{x}) = \min_{\vec{x} \in \mathbb{R}^n} F_0(\vec{x})$ (and the argmins are equal too). Hence the problem in (7.6) is equivalent to the original problem \mathcal{P} .

Thus through this reformulation, we have removed the explicit constraints of the problem and incorporated them into the objective function of the problem, and we can nominally write the problem as an unconstrained problem. If we had some magic algorithm which allowed us to solve *all* unconstrained problems, then this reduction would let us solve constrained problems as well.

Unfortunately, we do *not* have such a magic algorithm. Our usual algorithms for solving unconstrained problems, such as gradient descent, require differentiable objective functions that are well-defined. Thus, it falls to us to express our indicator functions in a differentiable way.

For this we consider *approximating* the indicator functions by other functions that are differentiable.

The key idea here is that, for a given indicator $\mathbb{1}[f_i(\vec{x}) \leq 0]$, we can express it as

$$\mathbb{1}[f_i(\vec{x}) \leq 0] = \max_{\lambda_i \in \mathbb{R}_+} \lambda_i f_i(\vec{x}), \quad (7.8)$$

where \mathbb{R}_+ is the set of non-negative real numbers. Why does this equality hold? Let's break it up into cases. Suppose that $f_i(\vec{x}) > 0$. Then λ_i being more positive would make $\lambda_i f_i(\vec{x})$ more positive, so the maximization will send $\lambda_i \rightarrow \infty$, making $\lambda_i f_i(\vec{x}) = \infty$. Now suppose that $f_i(\vec{x}) \leq 0$. Then λ_i being more positive would make $\lambda_i f_i(\vec{x})$ more negative, so the maximization will keep λ_i as its lower bound 0, making $\lambda_i f_i(\vec{x}) = 0$. For a similar reason, for an indicator $\mathbb{1}[h_j(\vec{x}) = 0]$, we have

$$\mathbb{1}[h_j(\vec{x}) = 0] = \max_{\nu_j \in \mathbb{R}} \nu_j h_j(\vec{x}). \quad (7.9)$$

Thus, we can rewrite the problem \mathcal{P} as

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \left\{ f_0(\vec{x}) + \sum_{i=1}^m \max_{\lambda_i \in \mathbb{R}_+} \lambda_i f_i(\vec{x}) + \sum_{j=1}^p \max_{\nu_j \in \mathbb{R}} \nu_j h_j(\vec{x}) \right\} \quad (7.10)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ f_0(\vec{x}) + \max_{\vec{\lambda} \in \mathbb{R}_+^m} \sum_{i=1}^m \lambda_i f_i(\vec{x}) + \max_{\vec{\nu} \in \mathbb{R}^p} \sum_{j=1}^p \nu_j h_j(\vec{x}) \right\} \quad (7.11)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{\nu} \in \mathbb{R}^p}} \left\{ f_0(\vec{x}) + \sum_{i=1}^m \lambda_i f_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}) \right\}. \quad (7.12)$$

This interior function is called the *Lagrangian*, and is much easier to work with than the original unconstrained problem with indicator functions. Thus we have made our constrained problem into a (mostly) unconstrained problem,¹ at the expense of adding an extra max. We will see how to cope with this extra difficulty shortly.

More formally, we have the following definition:

Definition 155 (Lagrangian)

The *Lagrangian* of problem \mathcal{P} in (7.1) is the function $L: \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p \rightarrow \mathbb{R}$ given by

$$L(\vec{x}, \vec{\lambda}, \vec{\nu}) \doteq f_0(\vec{x}) + \sum_{i=1}^m \lambda_i f_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}). \quad (7.13)$$

Additionally, $\lambda_i, \nu_j \in \mathbb{R}$ are called *Lagrange multipliers*.

The important intuition behind the Lagrange multipliers is that they are *penalties* for violating their corresponding constraint. In particular, just like how the indicator function $\mathbb{1}[f_i(\vec{x}) \leq 0]$ assigns an infinite penalty for violating the constraint $f_i(\vec{x}) \leq 0$, the term $\lambda_i f_i(\vec{x})$ assigns a penalty $\lambda_i f_i(\vec{x})$ for violating the constraint $f_i(\vec{x}) \leq 0$. One can show that $\lambda_i f_i(\vec{x}) \leq \mathbb{1}[f_i(\vec{x}) \leq 0]$, so the Lagrangian penalty is a smooth lower bound to the indicator penalty, which is something like a hard-threshold. We can do similar analysis for the equality constraints $h_j(\vec{x}) = 0$.

As derived before, the primal problem can be expressed in terms of the Lagrangian as

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{\nu} \in \mathbb{R}^p}} L(\vec{x}, \vec{\lambda}, \vec{\nu}). \quad (7.14)$$

We conclude with two very important properties of the Lagrangian.

Proposition 156

For every $\vec{x} \in \mathbb{R}^n$, the function $(\vec{\lambda}, \vec{\nu}) \mapsto L(\vec{x}, \vec{\lambda}, \vec{\nu})$ is an affine function, and hence a concave function. (We also say that L is affine (resp. concave) in $\vec{\lambda}$ and $\vec{\nu}$.)

Proof. The proof follows from the definition of the Lagrangian:

$$L(\vec{x}, \vec{\lambda}, \vec{\nu}) = f_0(\vec{x}) + \sum_{i=1}^m \lambda_i f_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}) \quad (7.15)$$

is affine in $\vec{\lambda}$ and $\vec{\nu}$. □

Example 157. Consider the optimization problem

$$\begin{aligned} p^* &= \min_{x \in \mathbb{R}} 3x^2 \\ \text{s.t.} \quad & 2x^3 \leq 8. \end{aligned} \quad (7.16)$$

This problem is not convex, but we can still find its Lagrangian as

$$L(x, \lambda) = 3x^2 + \lambda(2x^3 - 8). \quad (7.17)$$

¹Not quite — the max is actually over $\mathbb{R}_+^m \times \mathbb{R}^p$ which is technically a constrained set. But this non-negativity constraint is usually much easier to handle than arbitrary constraints, in part because the constraint set is convex.

7.2 Weak Duality

Our foray into duality starts by considering the primal problem:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{v} \in \mathbb{R}^p}} L(\vec{x}, \vec{\lambda}, \vec{v}). \quad (7.18)$$

The *dual problem* is obtained by swapping the min and max:

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{v} \in \mathbb{R}^p}} \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}). \quad (7.19)$$

In general, the quantities p^* and d^* *do not have to be equal!* And swapping the min and the max is the generic definition of “duality”.

Definition 158 (Dual Problem)

For a primal problem \mathcal{P} as described in (7.1), its *dual problem* \mathcal{D} is defined as

$$\begin{aligned} \text{problem } \mathcal{D}: \quad d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{v} \in \mathbb{R}^p}} g(\vec{\lambda}, \vec{v}) \quad (7.20) \\ \text{s.t. } \lambda_i \geq 0, \quad \forall i \in \{1, \dots, m\}. \end{aligned}$$

Here the function $g: \mathbb{R}_+^m \times \mathbb{R}^p \rightarrow \mathbb{R}$ is the *dual function* and defined as

$$g(\vec{\lambda}, \vec{v}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}). \quad (7.21)$$

Thus, $g(\vec{\lambda}, \vec{v})$ can be computed as an unconstrained optimization problem over \vec{x} , in particular $g(\vec{\lambda}, \vec{v})$ is the minimum value of $L(\vec{x}, \vec{\lambda}, \vec{v})$ over all \vec{x} .

There are several important properties of the dual problem and dual function, which we summarize below.

Proposition 159

The dual function g is a concave function of $(\vec{\lambda}, \vec{v})$, regardless of any properties of \mathcal{P} .

Proof. We already showed that $L(\vec{x}, \vec{\lambda}, \vec{v})$ is an affine (and thus concave) function of $\vec{\lambda}$ and \vec{v} . Thus the function

$$g(\vec{\lambda}, \vec{v}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}) \quad (7.22)$$

is a pointwise minimum of concave functions and is thus concave. \square

Corollary 160. *The dual problem \mathcal{D} is always a convex problem, no matter what the primal problem \mathcal{P} is.*

Note that the minimization of a convex function or the maximization of a concave function are both considered “convex” problems.

This means we have analytic and algorithmic ways to solve the dual problem \mathcal{D} . If we can connect the solutions to the dual problem \mathcal{D} to the primal problem \mathcal{P} , this means that we have reduced the process of solving \mathcal{P} to the process of solving a convex optimization problem, and this is something we know how to do. In the rest of this chapter, we discuss when this reduction is possible.

Example 161. Let us compute the dual of the optimization problem

$$\begin{aligned} p^* &= \min_{x \in \mathbb{R}} 5x^2 \\ \text{s.t. } & 3x \leq 5. \end{aligned} \quad (7.23)$$

The Lagrangian is

$$L(x, \lambda) = 5x^2 + \lambda(3x - 5). \quad (7.24)$$

The dual function is

$$g(\lambda) = \min_x L(x, \lambda) = \min_x \{5x^2 + \lambda(3x - 5)\}. \quad (7.25)$$

For each $\lambda > 0$, the function $L(x, \lambda) = 5x^2 + \lambda(3x - 5)$ is bounded below (in x), convex, and differentiable, so to minimize it we set its derivative to 0. The optimal x^* is a function of the Lagrange multiplier λ , so we write it as $x^*(\lambda)$. We have

$$0 = \nabla_x L(x^*(\lambda), \lambda) \quad (7.26)$$

$$= 10x^*(\lambda) + 3\lambda \quad (7.27)$$

$$\implies x^*(\lambda) = -\frac{3}{10}\lambda. \quad (7.28)$$

Thus we can plug this optimal point back in to get g :

$$g(\lambda) = L(x^*(\lambda), \lambda) \quad (7.29)$$

$$= 5(x^*(\lambda))^2 + \lambda(3x^*(\lambda) - 5) \quad (7.30)$$

$$= \frac{5 \cdot 9}{100}\lambda^2 + \lambda\left(-\frac{3 \cdot 3}{10}\lambda - 5\right) \quad (7.31)$$

$$= -\frac{9}{20}\lambda^2 - 5\lambda. \quad (7.32)$$

Thus the dual problem is

$$d^* = \max_{\lambda \in \mathbb{R}_+} \left(-\frac{9}{20}\lambda^2 - 5\lambda\right). \quad (7.33)$$

One may solve this problem directly to obtain $\lambda^* = 0$ and $d^* = 0$.

We also have some more bounds for the Lagrangian in terms of f_0 , g , p^* , and d^* .

Proposition 162

Let $\vec{x} \in \Omega$, let $\vec{\lambda} \in \mathbb{R}_+^m$, and let $\vec{v} \in \mathbb{R}^p$, so that \vec{x} is feasible for \mathcal{P} and $(\vec{\lambda}, \vec{v})$ is feasible for \mathcal{D} . Then we have:

- (a) $f_0(\vec{x}) \geq p^*$ and $g(\vec{\lambda}, \vec{v}) \leq d^*$;
- (b) $f_0(\vec{x}) \geq L(\vec{x}, \vec{\lambda}, \vec{v}) \geq g(\vec{\lambda}, \vec{v})$;
- (c) $f_0(\vec{x}) \geq d^*$ and $g(\vec{\lambda}, \vec{v}) \leq p^*$.

Proof. The inequalities in (a) follow from the characterizations

$$p^* = \min_{\vec{x}' \in \Omega} f_0(\vec{x}') \leq f_0(\vec{x}) \quad (7.34)$$

$$d^* = \max_{\substack{\vec{\lambda}' \in \mathbb{R}_+^m \\ \vec{v}' \in \mathbb{R}^p}} g(\vec{\lambda}', \vec{v}') \geq g(\vec{\lambda}, \vec{v}). \quad (7.35)$$

The inequalities in (b) follow from the characterizations

$$g(\vec{\lambda}, \vec{v}) = \min_{\vec{x}' \in \Omega} L(\vec{x}', \vec{\lambda}, \vec{v}) \leq L(\vec{x}, \vec{\lambda}, \vec{v}) \quad (7.36)$$

$$f_0(\vec{x}) = \max_{\substack{\vec{\lambda}' \in \mathbb{R}_+^m \\ \vec{v}' \in \mathbb{R}^p}} L(\vec{x}, \vec{\lambda}', \vec{v}') \geq L(\vec{x}, \vec{\lambda}, \vec{v}). \quad (7.37)$$

The inequalities in (c) follow from (b), in the sense that

$$f_0(\vec{x}) \geq g(\vec{\lambda}, \vec{v}) \quad \forall \vec{\lambda} \in \mathbb{R}_+^m \quad \forall \vec{v} \in \mathbb{R}^p \quad (7.38)$$

$$\implies f_0(\vec{x}) \geq \max_{\substack{\vec{\lambda}' \in \mathbb{R}_+^m \\ \vec{v}' \in \mathbb{R}^p}} g(\vec{\lambda}', \vec{v}') = d^*, \quad (7.39)$$

and additionally

$$f_0(\vec{x}) \geq g(\vec{\lambda}, \vec{v}) \quad \forall \vec{x} \in \Omega \quad (7.40)$$

$$\implies p^* = \min_{\vec{x}' \in \Omega} f_0(\vec{x}') \geq g(\vec{\lambda}, \vec{v}). \quad (7.41)$$

□

From all these inequalities, the relationship between p^* and d^* is totally unclear. Actually, their relationship is very simple; it is captured in a condition called “weak duality”.

Definition 163 (Types of Duality)

Let \mathcal{P} be a primal problem with optimum p^* . Let \mathcal{D} be the corresponding dual problem with optimum d^* .

- (a) If $p^* \geq d^*$ then we say that *weak duality* holds.
- (b) If $p^* = d^*$ then we say that *strong duality* holds.
- (c) The quantity $p^* - d^*$ is called the *duality gap*.

Proposition 164 (Weak Duality Always Holds)

For *any* problem, weak duality holds, i.e., the duality gap is non-negative.

The fact that weak duality always holds is actually a corollary of a slightly more generic result called the minimax inequality, which we now state and prove.

Proposition 165 (Minimax Inequality)

Let X and Y be *any* sets, and $F: X \times Y \rightarrow \mathbb{R}$ be *any* function. Then

$$\min_{x \in X} \max_{y \in Y} F(x, y) \geq \max_{y \in Y} \min_{x \in X} F(x, y). \quad (7.42)$$

Proof. For any $x \in X$ and $y \in Y$, we have

$$F(x, y) \geq \min_{x' \in X} F(x', y). \quad (7.43)$$

Taking the maximum over y on both sides (we will discuss why we can do this at the end of the proof), we have

$$\max_{y' \in Y} F(x, y') \geq \max_{y' \in Y} \min_{x' \in X} F(x', y'). \quad (7.44)$$

Finally, the left hand side is a function of x but the right hand side is a constant, so we might as well take the minimum in x on the left hand side to get

$$\min_{x' \in X} \max_{y' \in Y} F(x', y') \geq \max_{y' \in Y} \min_{x' \in X} F(x', y'). \quad (7.45)$$

This concludes the (main part of the) proof.

However, we still need to discuss why we can take maximums on both sides in Equation (7.44). We prove this in the case where the maximums are attained, though it is true in general (and can be proved via a slightly technical modification to this argument). Fix $x \in X$. Let $f: Y \rightarrow \mathbb{R}$ be defined as $f(y) \doteq F(x, y)$, and let $g: Y \rightarrow \mathbb{R}$ be defined as $g(y) \doteq \min_{x' \in X} F(x', y)$. Then we had shown that $f(y) \geq g(y)$ for all $y \in Y$, and we want to justify why this means that $\max_{y' \in Y} f(y') \geq \max_{y' \in Y} g(y')$. Towards this end, let $\tilde{y} \in \operatorname{argmax}_{y' \in Y} g(y')$. Then we have

$$\max_{y' \in Y} f(y') \geq f(\tilde{y}) \geq g(\tilde{y}) = \max_{y' \in Y} g(y') \implies \max_{y' \in Y} F(x, y') \geq \max_{y' \in Y} \min_{x' \in X} F(x', y') \text{ for all } x \in X, \quad (7.46)$$

as desired. \square

Here is a game theoretic interpretation of the minimax inequality.

Suppose X and Y are two players choosing actions x and y . Player X chooses action x trying to *minimize* the value of the function $F(x, y)$. Player Y chooses action y trying to *maximize* the value of the function $F(x, y)$. The two players are perfectly rational and take turns — that is, the players choose their actions one after the other, and the second player has full knowledge of the first player's action (which is locked in) as they choose their action. (In game theory this is called a two-player zero-sum sequential game.) The value of the game, given actions x and y , is $F(x, y)$.

The outer optimization corresponds to the player going first, and the inner minimization corresponds to the player going second. Why? Consider the optimization problem $\min_{x \in X} \max_{y \in Y} F(x, y)$. Then defining $G(x) = \max_{y \in Y} F(x, y)$, we see that given a particular x , $G(x)$ is the maximum value obtained by varying y of $F(x, y)$. That is, it corresponds to the best play by player Y given that player X picks action x . This means that the inner minimization corresponds to the second player Y 's actions given that the first player X chooses action x . And so the outer optimization corresponds to the first player X 's action. The situation is analogous on the right-hand side.

What the minimax inequality says is that *going second is better*. That is, being the second player and choosing your action with full knowledge of the first player's action leads to a better outcome for you than going first with no knowledge of the second player's action. When X is the second player, the value of the game is lower than when X is the first player; when Y is the second player, the value of the game is higher than when Y is the first player.

Given the minimax inequality, the proof that weak duality always holds is only one line.

Proof that weak duality always holds. We have

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{v} \in \mathbb{R}^p}} L(\vec{x}, \vec{\lambda}, \vec{v}) \geq \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{v} \in \mathbb{R}^p}} \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}) = d^*. \quad (7.47)$$

\square

Weak duality is useful because it gives a bound on how close to optimum a given decision variable is. More specifically, we have that for all $\vec{x} \in \Omega$ and all $(\vec{\lambda}, \vec{v}) \in \mathbb{R}_+^m \times \mathbb{R}^p$ that

$$f_0(\vec{x}) \geq p^* \geq d^* \geq g(\vec{\lambda}, \vec{v}) \implies f_0(\vec{x}) - g(\vec{\lambda}, \vec{v}) \geq f_0(\vec{x}) - p^*. \quad (7.48)$$

Thus, if $f_0(\vec{x}) - g(\vec{\lambda}, \vec{v}) \leq \epsilon$ then we know that $f_0(\vec{x}) - p^* \leq \epsilon$. This is called a *certificate of optimality*. We can use this certificate as a stopping condition for algorithms such as gradient descent or a broad class of algorithms known as *primal-dual* algorithms.

7.3 Strong Duality

Since weak duality holds for all problems, it necessarily would not have too many critical implications. The stronger condition of *strong duality* turns out to hold the keys to efficiently computing minima for constrained optimization problems.

First, a natural question is: when does strong duality hold? It does not *always* hold, and in particular it is a stronger condition than weak duality. You will see some example problems where strong duality does not hold in discussion.

There are many conditions under which strong duality holds. In this course, we discuss one such condition which is relatively simple to evaluate and broadly applicable. This condition is called *Slater's condition*.

Theorem 166 (Slater's Condition)

Consider a *convex* problem \mathcal{P} :

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) && (7.1) \\ \text{s.t.} \quad & f_i(\vec{x}) \leq 0, && \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0, && \forall j \in \{1, \dots, p\}. \end{aligned}$$

If there exists any $\vec{x} \in \text{relint}(\Omega)^a$ which is *strictly feasible*, i.e., such that all of the following hold:

- for all $i \in \{1, \dots, m\}$ such that f_i is an *affine* function, we have $f_i(\vec{x}) \leq 0$;
- for all $i \in \{1, \dots, m\}$ such that f_i is *not* an affine function, we have $f_i(\vec{x}) < 0$;
- and for all $j \in \{1, \dots, p\}$, we have $h_j(\vec{x}) = 0$;

then strong duality holds for \mathcal{P} and its dual \mathcal{D} , i.e., the duality gap is 0.

^aRecall that this notation means the *relative interior* of the feasible set Ω , defined formally in Definition 102. Since the formal definition of the relative interior is out of scope for this semester, you may just consider it to be the interior of Ω , i.e., points not on the boundary of Ω , and we will not give you problems where the distinction matters.

This result is sometimes called *refined Slater's condition* as it is a slight modification of another condition which is also called Slater's condition.

We will not prove this in class; a proof sketch is in [1]. It uses the separating hyperplane theorem we proved earlier. This result is actually one of the main payoffs of the separating hyperplane theorem that we will see in this course.

Again, observe that we *only* need a *single* point which fulfills the three conditions in order to declare that Slater's condition holds. Moreover, this point *need not be optimal* for the problem — any point at all which satisfies the conditions will do.

Slater's condition has a geometric interpretation: if there is a point in the *relative interior* of the feasible set which is *strictly feasible*, then strong duality holds. In problems we are able to visualize, this makes it relatively easy to determine whether Slater's condition holds.




Figure 7.1: Points in the convex set $\Omega = \{(x, y) \in \mathbb{R}^2 : -5 \leq x \leq 5, \frac{25}{4} \geq y \geq \frac{1}{4}x^2\}$. Blue points are feasible points which fulfill the conditions of Slater's condition (thus ensuring that Slater's condition holds for the problem $\min_{\vec{x} \in \Omega} f_0(\vec{x})$ provided that f_0 is convex). Red points are feasible points which do not fulfill the conditions of Slater's condition. Brown points are infeasible points.

Example 167. In this example, we consider the equality-constrained minimum-norm problem

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \|\vec{x}\|_2^2 & (7.49) \\ \text{s.t. } & A\vec{x} = \vec{y}. \end{aligned}$$

This problem only has equality constraints, and so it only has the Lagrange multiplier $\vec{\nu}$. The Lagrangian of this problem is

$$L(\vec{x}, \vec{\nu}) = \|\vec{x}\|_2^2 + \sum_{j=1}^p \nu_j (A\vec{x} - \vec{y})_j \quad (7.50)$$

$$= \|\vec{x}\|_2^2 + \vec{\nu}^\top (A\vec{x} - \vec{y}). \quad (7.51)$$

The dual function is

$$g(\vec{\nu}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\nu}) \quad (7.52)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ \|\vec{x}\|_2^2 + \vec{\nu}^\top (A\vec{x} - \vec{y}) \right\}. \quad (7.53)$$

The Lagrangian is convex in \vec{x} and bounded below, and so it is minimized wherever its gradient in \vec{x} is $\vec{0}$. The optimal \vec{x}^* is a function of $\vec{\nu}$, so we write it as $\vec{x}^*(\vec{\nu})$. We have

$$\vec{0} = \nabla_{\vec{x}} L(\vec{x}^*(\vec{\nu}), \vec{\nu}) \quad (7.54)$$

$$= 2\vec{x}^*(\vec{\nu}) + A^\top \vec{\nu} \quad (7.55)$$

$$\implies \vec{x}^*(\vec{\nu}) = -\frac{1}{2} A^\top \vec{\nu}. \quad (7.56)$$

Thus we can plug this optimal point back in to get g :

$$g(\vec{\nu}) = L(\vec{x}^*(\vec{\nu}), \vec{\nu}) \quad (7.57)$$

$$= \|\vec{x}^*(\vec{\nu})\|_2^2 + \vec{\nu}^\top (A\vec{x}^*(\vec{\nu}) - \vec{y}) \quad (7.58)$$

$$= \left\| -\frac{1}{2}A^\top \vec{v} \right\|_2^2 + \vec{v}^\top \left(A \left(-\frac{1}{2}A^\top \vec{v} \right) - \vec{y} \right) \quad (7.59)$$

$$= \frac{1}{4}\vec{v}^\top AA^\top \vec{v} - \frac{1}{2}\vec{v}^\top AA^\top \vec{v} - \vec{v}^\top \vec{y} \quad (7.60)$$

$$= -\frac{1}{4}\vec{v}^\top AA^\top \vec{v} - \vec{v}^\top \vec{y}. \quad (7.61)$$

Thus the dual problem is

$$d^* = \max_{\vec{v} \in \mathbb{R}^p} \left\{ -\frac{1}{4}\vec{v}^\top AA^\top \vec{v} - \vec{v}^\top \vec{y} \right\} = - \min_{\vec{v} \in \mathbb{R}^p} \left\{ \frac{1}{4}\vec{v}^\top AA^\top \vec{v} + \vec{v}^\top \vec{y} \right\}. \quad (7.62)$$

There are no constraints on the dual problem, because there are no inequality constraints on the primal problem. Since this dual problem is a convex unconstrained problem whose objective value is bounded below, it can be solved by setting the gradient of the objective to $\vec{0}$ and solving for the optimal dual variable. This yields

$$\vec{0} = \nabla g(\vec{v}^*) \quad (7.63)$$

$$= \frac{1}{2}AA^\top \vec{v}^* + \vec{y} \quad (7.64)$$

$$\implies \vec{v}^* = -2(AA^\top)^{-1}\vec{y}. \quad (7.65)$$

Our original primal problem is convex, and all constraints are affine. As long as there is a feasible point \vec{x} , i.e., a point \vec{x} such that $A\vec{x} = \vec{y}$, then Slater's condition implies that strong duality holds. Suppose that there is such a feasible point (i.e., the primal problem is feasible). Then strong duality holds. Because the primal problem is convex and strong duality holds, we can recover the optimal primal variable \vec{x}^* from the optimal dual variable \vec{v}^* as

$$\vec{x}^* = \vec{x}^*(\vec{v}^*) = -\frac{1}{2}A^\top (-2(AA^\top)^{-1}\vec{y}) = A^\top (AA^\top)^{-1}\vec{y}. \quad (7.66)$$

This corresponds to the familiar minimum-norm solution.

Example 168. Consider a so-called *linear program*:

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \\ \text{s.t.} \quad & A\vec{x} = \vec{y} \\ & \vec{x} \geq \vec{0} \end{aligned} \quad (7.67)$$

where the last constraint means $x_i \geq 0$ for all i . Slater's condition implies that strong duality holds for this problem provided there is a feasible point. The Lagrangian is

$$L(\vec{x}, \vec{\lambda}, \vec{v}) = \vec{c}^\top \vec{x} + \sum_{i=1}^n \lambda_i (-x_i) + \sum_{j=1}^m \nu_j (A\vec{x} - \vec{y})_j \quad (7.68)$$

$$= \vec{c}^\top \vec{x} - \vec{\lambda}^\top \vec{x} + \vec{v}^\top (A\vec{x} - \vec{y}) \quad (7.69)$$

$$= \left(\vec{c} + A^\top \vec{v} - \vec{\lambda} \right)^\top \vec{x} - \vec{v}^\top \vec{y}. \quad (7.70)$$

This function is affine in \vec{x} , hence convex in \vec{x} , so its dual function is

$$g(\vec{\lambda}, \vec{v}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}) \quad (7.71)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ \left(\vec{c} + A^\top \vec{v} - \vec{\lambda} \right)^\top \vec{x} - \vec{v}^\top \vec{y} \right\} \quad (7.72)$$

$$= \begin{cases} -\vec{v}^\top \vec{y}, & \text{if } \vec{c} + A^\top \vec{v} - \vec{\lambda} = \vec{0} \\ -\infty, & \text{otherwise.} \end{cases} \quad (7.73)$$

(It is a good exercise to figure out why this last equality is correct.) Thus, the dual is

$$\begin{aligned} d^* &= \max_{\substack{\vec{\lambda} \in \mathbb{R}^n \\ \vec{v} \in \mathbb{R}^m}} g(\vec{\lambda}, \vec{v}) \\ \text{s.t. } &\lambda_i \geq 0, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (7.74)$$

Expanding out g , this problem is equivalent to

$$\begin{aligned} d^* &= \max_{\substack{\vec{\lambda} \in \mathbb{R}^n \\ \vec{v} \in \mathbb{R}^m}} -\vec{v}^\top \vec{y} \\ \text{s.t. } &\lambda_i \geq 0, \quad \forall i \in \{1, \dots, n\} \\ &\vec{c} + A^\top \vec{v} - \vec{\lambda} = \vec{0}. \end{aligned} \quad (7.75)$$

Example 169 (Shadow Prices). In this example, we determine an economic interpretation of Lagrange multipliers.

Suppose we have 200 kilos of merlot grapes and 300 kilos of shiraz grapes. Consider the following possible blends:

- 4 kilos merlot, 1 kilo shiraz for \$20 per bottle.
- 2 kilos merlot, 3 kilos shiraz for \$15 per bottle.

We want to maximize our profits. Suppose we make q_1 bottles of the first blend, and q_2 of the second. Our optimization problem is:

$$\begin{aligned} p^* &= \max_{q_1, q_2 \in \mathbb{R}} 20q_1 + 15q_2 \\ \text{s.t. } &4q_1 + 2q_2 \leq 200 \\ &q_1 + 3q_2 \leq 300 \\ &q_1 \geq 0 \\ &q_2 \geq 0. \end{aligned} \quad (7.76)$$

In reality, we can't make fractional bottles of wine, so we can round q_1 and q_2 to the nearest integer if needed. Now consider the following modification. Suppose that we actually want to sell off some of the grapes instead of turning them into wine. We can earn λ_1 dollars per kilo of merlot and λ_2 per kilo of shiraz. This yields a new optimization problem

$$\max_{q_1, q_2 \in \mathbb{R}_+} \{20q_1 + 15q_2 + \lambda_1(200 - 4q_1 - 2q_2) + \lambda_2(300 - q_1 - 3q_2)\} \quad (7.77)$$

$$= \max_{q_1, q_2 \in \mathbb{R}_+} \{(20 - 4\lambda_1 - \lambda_2)q_1 + (15 - 2\lambda_1 - 3\lambda_2)q_2 + 200\lambda_1 + 300\lambda_2\}. \quad (7.78)$$

If the coefficient of q_1 is negative, we shouldn't make any of the first blend. If the coefficient of q_2 is negative, we shouldn't make any of the second blend. If both are positive, we should make as many of each as possible. What happens when $20 - 4\lambda_1 - \lambda_2 = 0$ and $15 - 2\lambda_1 - 3\lambda_2 = 0$? This is an indifference point, i.e. the point at which our profit is the same no matter how many bottles of either wine we make. Under this condition, the minimum profit we could possibly make is

$$\min_{\lambda_1, \lambda_2 \in \mathbb{R}_+} 200\lambda_1 + 300\lambda_2 \quad (7.79)$$

$$\begin{aligned} \text{s.t. } 20 - 4\lambda_1 - \lambda_2 &= 0 \\ 15 - 2\lambda_1 - 3\lambda_2 &= 0. \end{aligned}$$

This problem is the *dual* to our original (primal) problem. The λ_i 's are called *shadow prices*, and capture how much we're willing to pay to violate our constraints.

7.4 Karush-Kuhn-Tucker (KKT) Conditions

We now go into some more involved theory which connects strong duality to optimality conditions.

A broadly applicable set of conditions which are sometimes necessary and sometimes sufficient for optimality is called the *Karush-Kuhn-Tucker* (KKT) conditions.

Definition 170 (KKT Conditions)

Let $(\vec{x}, \vec{\lambda}, \vec{\nu}) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p$ be a decision variable and Lagrange multipliers. Suppose that the objective function f_0 and constraint functions $f_1, \dots, f_m, h_1, \dots, h_p$ are differentiable. We say that $(\vec{x}, \vec{\lambda}, \vec{\nu})$ fulfills the *KKT conditions* if

- (Primal feasibility.) \vec{x} is feasible for \mathcal{P} , i.e.,

$$f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\} \quad (7.80)$$

$$\text{and } h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}. \quad (7.81)$$

- (Dual feasibility.) $(\vec{\lambda}, \vec{\nu})$ is feasible for \mathcal{D} , i.e.,

$$\tilde{\lambda}_i \geq 0, \quad \forall i \in \{1, \dots, m\}. \quad (7.82)$$

- (Complementary slackness.)

$$\tilde{\lambda}_i f_i(\vec{x}) = 0, \quad \forall i \in \{1, \dots, m\}. \quad (7.83)$$

- (Stationarity, or first-order condition.)

$$\vec{0} = \nabla_{\vec{x}} L(\vec{x}, \vec{\lambda}, \vec{\nu}) \quad (7.84)$$

$$= \nabla f(\vec{x}) + \sum_{i=1}^m \tilde{\lambda}_i \nabla f_i(\vec{x}) + \sum_{j=1}^p \tilde{\nu}_j \nabla h_j(\vec{x}). \quad (7.85)$$

Given an arbitrary optimization problem, the KKT conditions do *not* have to be related to the optimality conditions. But it turns out that in many cases, they are related.

Theorem 171 (If Strong Duality Holds, then KKT Conditions are Necessary for Optimality)

Suppose \mathcal{P} is a primal problem with dual \mathcal{D} . Suppose that the objective function f_0 and constraint functions $f_1, \dots, f_m, h_1, \dots, h_p$ are differentiable, strong duality holds, and $(\vec{x}^*, \vec{\lambda}^*, \vec{\nu}^*)$ are optimal primal and dual variables. Then $(\vec{x}^*, \vec{\lambda}^*, \vec{\nu}^*)$ fulfill the KKT conditions.

Proof. By assumption, \vec{x}^* is feasible for \mathcal{P} and $(\vec{\lambda}^*, \vec{\nu}^*)$ is feasible for \mathcal{D} . This implies that $\lambda_i^* f_i(\vec{x}^*) \leq 0$ for all i , and

$\nu_j^* h_j(\bar{x}^*) = 0$ for all j . Thus, we have

$$d^* = g(\bar{\lambda}^*, \bar{\nu}^*) \quad (7.86)$$

$$= \min_{\bar{x} \in \mathbb{R}^n} L(\bar{x}, \bar{\lambda}^*, \bar{\nu}^*) \quad (7.87)$$

$$\leq L(\bar{x}^*, \bar{\lambda}^*, \bar{\nu}^*) \quad (7.88)$$

$$= f_0(\bar{x}^*) + \underbrace{\sum_{i=1}^m \lambda_i^* f_i(\bar{x}^*)}_{\leq 0} + \underbrace{\sum_{j=1}^p \nu_j^* h_j(\bar{x}^*)}_{=0} \quad (7.89)$$

$$\leq f_0(\bar{x}^*) \quad (7.90)$$

$$= p^*. \quad (7.91)$$

Because $p^* = d^*$, all inequalities in the above chain are actually equalities. This means that \bar{x}^* minimizes $L(\bar{x}, \bar{\lambda}^*, \bar{\nu}^*)$ over $\bar{x} \in \mathbb{R}^n$. By the main theorem of vector calculus, this implies that $\nabla_{\bar{x}} L(\bar{x}^*, \bar{\lambda}^*, \bar{\nu}^*) = \vec{0}$, which is the stationarity condition. It also implies that $\sum_{i=1}^m \lambda_i^* f_i(\bar{x}^*) = 0$. But each term in the sum is ≤ 0 , so they must all be 0. This means that $\lambda_i^* f_i(\bar{x}^*) = 0$ for each i . This gives the complementary slackness condition. Thus the KKT conditions hold for $(\bar{x}^*, \bar{\lambda}^*, \bar{\nu}^*)$. \square

Theorem 172 (If Convexity Holds, then KKT Conditions are Sufficient for Optimality)

Suppose \mathcal{P} is a primal problem with dual \mathcal{D} . Suppose that the objective function f_0 and constraint functions $f_1, \dots, f_m, h_1, \dots, h_p$ of \mathcal{P} are differentiable. Suppose that f_0, f_1, \dots, f_m are convex and h_1, \dots, h_p are affine. Suppose that $(\vec{x}, \vec{\lambda}, \vec{\nu})$ fulfill the KKT conditions. Then \mathcal{P} is convex, strong duality holds, and $(\vec{x}, \vec{\lambda}, \vec{\nu})$ are optimal primal and dual variables.

Proof. By assumption, \vec{x} is feasible for \mathcal{P} and $(\vec{\lambda}, \vec{\nu})$ is feasible for \mathcal{D} . Since the f_i are convex and the h_j are affine, \mathcal{P} is convex, and

$$L(\vec{x}, \vec{\lambda}, \vec{\nu}) = f_0(\vec{x}) + \sum_{i=1}^m \lambda_i f_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}) \quad (7.92)$$

is convex in \vec{x} . Since stationarity holds, we have

$$\nabla_{\vec{x}} L(\vec{x}, \vec{\lambda}, \vec{\nu}) = \vec{0}, \quad (7.93)$$

so because the primal problem is convex, we have that \vec{x} minimizes $L(\vec{x}, \vec{\lambda}, \vec{\nu})$ over all $\vec{x} \in \mathbb{R}^n$. Thus

$$g(\vec{\lambda}, \vec{\nu}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{\nu}) \quad (7.94)$$

$$= L(\vec{x}, \vec{\lambda}, \vec{\nu}) \quad (7.95)$$

$$= f_0(\vec{x}) + \underbrace{\sum_{i=1}^m \tilde{\lambda}_i f_i(\vec{x})}_{=0} + \underbrace{\sum_{j=1}^p \tilde{\nu}_j h_j(\vec{x})}_{=0} \quad (7.96)$$

$$= f_0(\vec{x}). \quad (7.97)$$

Thus $f_0(\vec{x}) = g(\vec{\lambda}, \vec{\nu})$, so the duality gap is 0 and strong duality holds, with $(\vec{x}, \vec{\lambda}, \vec{\nu})$ being optimal primal and dual variables. \square

In this course, most problems will be convex and strong duality will hold; in this case the above two theorems can apply and we get that the KKT conditions are *equivalent* to optimality conditions. This is the source of the intuition that convex problems are easier to optimize.

Corollary 173 (If Convexity and Strong Duality Hold, then KKT Conditions are Necessary and Sufficient for Optimality). *Suppose \mathcal{P} is a primal problem with dual \mathcal{D} . Suppose strong duality holds for \mathcal{P} and \mathcal{D} . Suppose that the objective function f_0 and constraint functions $f_1, \dots, f_m, h_1, \dots, h_p$ for \mathcal{P} are differentiable. Suppose that f_0, f_1, \dots, f_m are convex and h_1, \dots, h_p are affine. Let $(\vec{x}, \vec{\lambda}, \vec{v}) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p$ be primal and dual variables. Then $(\vec{x}, \vec{\lambda}, \vec{v})$ are optimal primal and dual variables if and only if they fulfill the KKT conditions.*

The following is a generic sequence of steps that you can try for any convex optimization problem.

Problem Solving Strategy 174 (Solving Convex Optimization Problems Using KKT Conditions). *Suppose you have a problem \mathcal{P} with dual \mathcal{D} .*

1. *Show that \mathcal{P} is convex and the objective and constraint functions are differentiable.*
2. *Show Slater's condition holds and/or that strong duality holds for \mathcal{P} and \mathcal{D} .*
3. *Compute the KKT conditions for \mathcal{P} and \mathcal{D} .*
4. *Solve for the optimal primal and dual variables using the KKT conditions.*

Even for more complicated problems where it is not possible to solve the KKT conditions analytically, many algorithms can be interpreted as solving the KKT conditions under various conditions.

Example 175 (Example 161, with KKT Conditions). In this example, we apply the KKT conditions to the problem in Example 161. Consider the following problem:

$$p^* = \min_{x \in \mathbb{R}} 5x^2 \tag{7.98}$$

$$\text{s.t. } 3x \leq 5. \tag{7.99}$$

Its Lagrangian is

$$L(x, \lambda) = 5x^2 + \lambda(3x - 5). \tag{7.100}$$

The problem is convex, and there exists a strictly feasible point in the relative interior of the feasible set, e.g., $x = 0$ (notice that this is also global optimum, but it does not have to be; $x = -1$ would have sufficed just as well to be a strictly feasible point in the relative interior of the feasible set). Thus Slater's condition holds, strong duality holds, and the KKT conditions are necessary and sufficient for optimality. Now we solve the system to global optimum using KKT conditions.

Let $(\tilde{x}, \tilde{\lambda})$ solve the KKT conditions. Then they must obey:

- Primal feasibility: $3\tilde{x} \leq 5$.
- Dual feasibility: $\tilde{\lambda} \geq 0$.
- Complementary slackness: $\tilde{\lambda}(3\tilde{x} - 5) = 0$.
- Stationarity: $0 = \nabla_x L(\tilde{x}, \tilde{\lambda}) = 10\tilde{x} + 3\tilde{\lambda}$.

Now we solve for $\tilde{\lambda}$ in terms of \tilde{x} . We obtain from stationarity that

$$\tilde{\lambda} = -\frac{10}{3}\tilde{x}. \quad (7.101)$$

Thus by complementary slackness we have

$$0 = -\frac{10}{3}\tilde{x}(3\tilde{x} - 5), \quad (7.102)$$

so that

$$\tilde{x} = 0 \quad \text{or} \quad \tilde{x} = \frac{5}{3}. \quad (7.103)$$

Supposing that $\tilde{x} = 5/3$, then we have $\tilde{\lambda} = -5/3$, which violates dual feasibility. Thus we must have $\tilde{x} = 0$, implying $\tilde{\lambda} = 0$. And this indeed satisfies all of the KKT conditions. In particular, we can tell from inspection that at least the primal variable $\tilde{x} = 0 = x^*$ is globally optimal.

Note that even when solving the KKT system correctly, we (momentarily) ended up with an answer that would not satisfy the KKT conditions in the first place. This is one way to solve such systems: simplify the system as much as possible, generate two or several possible solutions $(\tilde{x}, \tilde{\lambda})$, and check again to see which one(s) still make the KKT conditions hold.

The following content is optional/out of scope for this semester. Regardless, it may be helpful to read it to gain context, or get a deeper understanding of various results.

7.5 (OPTIONAL) Conic Duality

Below, we introduce how Lagrange duality can be extended to optimization problems with inequality constraints defined via generalized inequalities, as defined below. For more details, please see [1, Section 5.9].

Definition 176 (Generalized Inequalities using Convex Cones)

Let $K \subseteq \mathbb{R}^n$ be a proper cone. Let $\text{int}(K)$ be the interior of K . The *generalized inequality* induced by K is a relation \succeq_K defined on pairs of vectors $\vec{u}, \vec{v} \in \mathbb{R}^n$ with the following properties:

- (a) If $\vec{u} - \vec{v} \in K$, we write $\vec{u} \succeq_K \vec{v}$ and $\vec{v} \preceq_K \vec{u}$.
- (b) If $\vec{u} - \vec{v} \in \text{int}(K)$, we write $\vec{u} \succ_K \vec{v}$ and $\vec{v} \prec_K \vec{u}$.

Substituting $\vec{0}$ into the above notation, we get the following notations:

- (c) If $\vec{u} \in K$, we write $\vec{u} \succeq_K \vec{0}$.
- (d) If $-\vec{v} \in K$, we write $\vec{v} \preceq_K \vec{0}$.
- (e) If $\vec{u} \in \text{int}(K)$, we write $\vec{u} \succ_K \vec{0}$.
- (f) If $-\vec{v} \in \text{int}(K)$, we write $\vec{v} \prec_K \vec{0}$.

It is possible that $\vec{u} - \vec{v} \notin K$, in which case there is no generalized inequality with respect to K between them.

Example 177. While the above may look scary, we have already seen generalized inequalities in this course. The set of positive definite matrices \mathbb{S}_+^n is a proper cone in the set of symmetric matrices \mathbb{S}^n . Thus when we write that $A \succeq 0$ for a symmetric matrix A , we are really using the generalized inequality with respect to the cone $K = \mathbb{S}_+^n$. Correspondingly, we can write $A \succ B$ to mean $A - B$ is positive definite.

Generalized inequalities can also help simplify or generalize several concepts introduced earlier in the course. Suppose for example that we have a familiar convex optimization problem:

$$\min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (7.104)$$

$$\text{s.t. } f_i(\vec{x}) \leq 0, \quad \forall i \in \{1, \dots, m\}. \quad (7.105)$$

(One can also add equality constraints to this problem without issue.) Now consider the set $K = \mathbb{R}_+^m = \{\vec{x} \in \mathbb{R}^m : x_i \geq 0 \forall i \in \{1, \dots, m\}\}$, i.e., the *non-negative orthant*. Indeed, K is a proper cone (proof is an exercise). If we collect all constraints into a single vector-valued constraint function $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ given by

$$\vec{f}(\vec{x}) = \begin{bmatrix} f_1(\vec{x}) \\ \vdots \\ f_m(\vec{x}) \end{bmatrix} \quad (7.106)$$

then the constraint $\vec{f}(\vec{x}) \preceq_K \vec{0}$ means that $-\vec{f}(\vec{x}) \in K$, that is, each $-f_i(\vec{x}) \geq 0$, or equivalently $f_i(\vec{x}) \leq 0$. These forms of the m constraints are absolutely equivalent. Thus, our original familiar convex optimization problem is equivalent to the following problem:

$$\min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (7.107)$$

$$\text{s.t. } \vec{f}(\vec{x}) \preceq_K \vec{0}. \quad (7.108)$$

We will work with generalizations of this problem shortly.

Observe that, for any proper cone $K \subseteq \mathbb{R}^n$ and given $\vec{v}_1, \vec{v}_2, \vec{v}_3 \in \mathbb{R}^n$, if we have $\vec{v}_1 \succeq_K \vec{v}_2$ and $\vec{v}_2 \succeq_K \vec{v}_3$, then $\vec{v}_1 \succeq_K \vec{v}_3$.

Now we use this generalized inequality system to define more general convex optimization problems. Let $f_0: \mathbb{R}^n \rightarrow \mathbb{R}$ be the objective function. For each $i \in \{1, \dots, m\}$, let $\vec{f}_i: \mathbb{R}^n \rightarrow \mathbb{R}^{d_i}$ be a *vector-valued* inequality constraint function (we will see shortly how this works), and let $K_i \subseteq \mathbb{R}^{d_i}$ be a proper cone (this could be called the *constraint cone*). For convenience, let $d \doteq \sum_{i=1}^m d_i$. Finally, for each $j \in \{1, \dots, p\}$, let $h_j: \mathbb{R}^n \rightarrow \mathbb{R}$ be a (usual scalar-valued) equality constrained function.

Consider the following primal optimization problem over \mathbb{R}^n with generalized inequality constraints:

$$\text{problem } \mathcal{P}: \quad p^* = \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (7.109)$$

$$\text{s.t. } \vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0}, \quad \forall i \in \{1, \dots, m\}$$

$$h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}.$$

We aim to derive a theory of generalized Lagrangian duality for this system.

As before, let $\mathbb{1}[\cdot]$ be an indicator function that returns 0 when the input condition is true, and $+\infty$ otherwise. Then, as before, we can look at the unconstrained variant:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \left[f_0(\vec{x}) + \sum_{i=1}^m \mathbb{1}[\vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0}] + \sum_{j=1}^p \mathbb{1}[h_j(\vec{x}) = 0] \right]. \quad (7.110)$$

Recall that we showed, en route to the Lagrangian duality presented in Section 7.1, that

$$\mathbb{1}[f_i(\vec{x}) \leq 0] = \max_{\lambda_i \in \mathbb{R}_+} \lambda_i f_i(\vec{x}), \quad \mathbb{1}[h_j(\vec{x}) = 0] = \max_{\nu_j \in \mathbb{R}} \nu_j h_j(\vec{x}). \quad (7.111)$$

The latter equality will help us here again, but the former will not, and we have to derive an analogue. Indeed, we claim that

$$\mathbb{1} \left[\vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0} \right] = \max_{\substack{\vec{\lambda}_i \in \mathbb{R}^{d_i} \\ \vec{\lambda}_i \succeq_{K_i^*} \vec{0}}} \vec{\lambda}_i^\top \vec{f}_i(\vec{x}), \quad (7.112)$$

where K_i^* denotes the dual cone of K_i in \mathbb{R}^{d_i} . Why is this true? It turns out to be for a similar reason as the more special case above. If $\vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0}$, then $-\vec{f}_i(\vec{x}) \in K_i$. By definition, for $\vec{\lambda}_i \in K_i^*$, we must have

$$0 \leq \vec{\lambda}_i^\top (-\vec{f}_i(\vec{x})) = -\vec{\lambda}_i^\top \vec{f}_i(\vec{x}) \implies \vec{\lambda}_i^\top \vec{f}_i(\vec{x}) \leq 0. \quad (7.113)$$

Equality in this case is obtained by selecting $\vec{\lambda}_i = \vec{0}$, which we are assured is legal because, by definition of a proper cone, $\vec{0} \in K_i$. Thus we have shown

$$\vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0} \implies \max_{\substack{\vec{\lambda}_i \in \mathbb{R}^{d_i} \\ \vec{\lambda}_i \succeq_{K_i^*} \vec{0}}} \vec{\lambda}_i^\top \vec{f}_i(\vec{x}) = 0. \quad (7.114)$$

In the other case, suppose $\vec{f}_i(\vec{x}) \not\preceq_{K_i} \vec{0}$. Then we have $-\vec{f}_i(\vec{x}) \notin K_i$. Since K_i is a proper cone, it is closed and convex, so $K_i = (K_i^*)^*$. We thus have $-\vec{f}_i(\vec{x}) \notin (K_i^*)^*$. Namely, there exists some $\vec{\lambda}_i \in K_i^*$ such that $\vec{\lambda}_i^\top (-\vec{f}_i(\vec{x})) < 0$, or equivalently $\vec{\lambda}_i^\top \vec{f}_i(\vec{x}) > 0$. Since K_i^* is a cone, it is closed under positive scalar multiplication; thus, for any $\alpha_i > 0$, we have $\alpha_i \vec{\lambda}_i \in K_i^*$, so

$$(\alpha_i \vec{\lambda}_i)^\top \vec{f}_i(\vec{x}) = \alpha_i (\vec{\lambda}_i^\top \vec{f}_i(\vec{x})) > 0 \quad (7.115)$$

Taking $\alpha_i \rightarrow \infty$ obtains that the maximum over $\vec{\lambda}_i$ is $+\infty$ in this case. Namely, we have shown

$$\vec{f}_i(\vec{x}) \not\preceq_{K_i} \vec{0} \implies \max_{\substack{\vec{\lambda}_i \in \mathbb{R}^{d_i} \\ \vec{\lambda}_i \succeq_{K_i^*} \vec{0}}} \vec{\lambda}_i^\top \vec{f}_i(\vec{x}) = \infty. \quad (7.116)$$

Thus, defining $\vec{\lambda} = (\vec{\lambda}_1, \dots, \vec{\lambda}_m) \in \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_m} = \mathbb{R}^d$ for convenience, we have

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \max_{\substack{\vec{\lambda} \in \mathbb{R}^d \\ \vec{v} \in \mathbb{R}^p \\ \vec{\lambda}_i \preceq_{K_i} \vec{0} \forall i \in \{1, \dots, m\}}} \left[f_0(\vec{x}) + \sum_{i=1}^m \vec{\lambda}_i^\top \vec{f}_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}) \right]. \quad (7.117)$$

The above discussion inspires the following definition of a generalized Lagrangian $L: \mathbb{R}^n \times \mathbb{R}^d \times \mathbb{R}^p$ for the given primal optimization problem with generalized inequalities.²

$$L(\vec{x}, \vec{\lambda}, \vec{v}) \doteq f_0(\vec{x}) + \sum_{i=1}^m \vec{\lambda}_i^\top \vec{f}_i(\vec{x}) + \sum_{j=1}^p \nu_j h_j(\vec{x}). \quad (7.118)$$

As before, we define the dual function $g: \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ by

$$g(\vec{\lambda}, \vec{v}) \doteq \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{v}). \quad (7.119)$$

Thus, we can write the dual problem of (7.109) as

$$\text{problem } \mathcal{D}: \quad d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^d \\ \vec{v} \in \mathbb{R}^p}} g(\vec{\lambda}, \vec{v}) \quad (7.120)$$

²If the optimization problem under study were over a general inner product space, then the expressions $\vec{\lambda}_i^\top \vec{f}_i(\vec{x})$ should be replaced with $\langle \vec{\lambda}_i, \vec{f}_i(\vec{x}) \rangle$. This occurs primarily in semidefinite programming, where \vec{f}_i could produce a symmetric matrix as an output. In this case, $\vec{\lambda}_i$ would similarly be a symmetric matrix, and the inner product would be the Frobenius inner product.

$$\text{s.t. } \vec{\lambda}_i \succeq_{K_i^*} \vec{0}, \quad \forall i \in \{1, \dots, m\}.$$

Note that weak duality, i.e., $d^* \leq p^*$, always holds, by the minimax inequality.

An analog of Slater's condition holds for optimization problems with generalized inequalities induced by proper cones. We first require the following definition, which extends the notion of the convexity of a function to generalized convexity, where the inequalities in the definition of convexity are turned into general inequalities induced by proper cones.

Definition 178 (Generalized Convexity)

Let $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^d$, and let $K \subseteq \mathbb{R}^d$ be a proper cone.

(a) We say that \vec{f} is K -convex if for any $\alpha \in [0, 1]$ and $\vec{x}, \vec{y} \in \mathbb{R}^n$, we have

$$\vec{f}(\alpha\vec{x} + (1 - \alpha)\vec{y}) \preceq_K \alpha\vec{f}(\vec{x}) + (1 - \alpha)\vec{f}(\vec{y}). \quad (7.121)$$

(b) We say that \vec{f} is strictly K -convex if for any $\alpha \in (0, 1)$ and $\vec{x}, \vec{y} \in \mathbb{R}^n$ with $\vec{x} \neq \vec{y}$, we have

$$\vec{f}(\alpha\vec{x} + (1 - \alpha)\vec{y}) \prec_K \alpha\vec{f}(\vec{x}) + (1 - \alpha)\vec{f}(\vec{y}). \quad (7.122)$$

Theorem 179 (Generalized Slater's Condition)

Consider a primal problem \mathcal{P} :

$$\begin{aligned} \text{problem } \mathcal{P}: \quad p^* &= \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) & (7.109) \\ \text{s.t.} \quad & \vec{f}_i(\vec{x}) \preceq_{K_i} \vec{0}, \quad \forall i \in \{1, \dots, m\} \\ & h_j(\vec{x}) = 0, \quad \forall j \in \{1, \dots, p\}, \end{aligned}$$

where

- the function $f_0: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex.;
- for each $i \in \{1, \dots, m\}$, the function $\vec{f}_i: \mathbb{R}^n \rightarrow \mathbb{R}^{d_i}$ is K_i -convex, where $K_i \subseteq \mathbb{R}^{d_i}$ is a proper cone;
- for each $j \in \{1, \dots, p\}$, the function $h_j: \mathbb{R}^n \rightarrow \mathbb{R}$ is affine.

If there exists any point $\vec{x} \in \text{relint}(\Omega)$ which is strictly feasible, i.e., such that all of the following holds:

- for each $i \in \{1, \dots, m\}$, we have $\vec{f}_i(\vec{x}) \prec_{K_i} \vec{0}$;
- and for each $j \in \{1, \dots, p\}$, we have $h_j(\vec{x}) = 0$;

then strong duality holds for \mathcal{P} and its dual \mathcal{D} , i.e., the duality gap is 0. ^a

^aAnalogous conditions hold for the general case where the primal optimization problem is defined over a vector space equipped with an inner product and a norm, as in the case of semidefinite programming.

Chapter 8

Types of Optimization Problems

Relevant sections of the textbooks:

- [1] Chapter 4.
- [2] Chapters 9, 10, 12.

8.1 Linear Programs

In this chapter we introduce a *taxonomy* of common optimization problems that can be efficiently solved through a variety of ways. We use the notation $\vec{u} \geq \vec{v}$ to denote $u_i \geq v_i$ for all i .

A linear program is one with an affine objective function and affine constraint functions (both inequality and equality constraints). It is the most restrictive (e.g. smallest) class in the taxonomy we present. It has a *standard form* denoted below.

Definition 180 (Linear Program)

A *linear program* (LP) is an optimization problem with an affine objective and affine constraints. A *standard form linear program* is an optimization problem of the following form:^a

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} & (8.1) \\ \text{s.t. } & A\vec{x} = \vec{y} \\ & \vec{x} \geq \vec{0}. \end{aligned}$$

^aConstant terms have been omitted from the objective function.

There are many equivalent forms of a linear program; in particular, the following proposition (whose proof we leave as an exercise) can be shown using slack variables.

Proposition 181

Any linear program is equivalent to a standard form linear program.

Putting linear programs in standard form is important because the standard form is commonly accepted by optimization algorithms and implementations. Usually if you provide a linear program that isn't in standard form to a

solver, they will convert it to standard form first before running their algorithms. Conversions to-and-from standard form may increase the number of variables in the problem and the eventual algorithmic complexity of solving it. One example of this conversion is done below.

Example 182. Suppose we have the linear program

$$\min_{\vec{x} \in \mathbb{R}^2} 2x_1 + 4x_2 \quad (8.2)$$

$$\text{s.t. } x_1 + x_2 \geq 3 \quad (8.3)$$

$$3x_1 + 2x_2 = 14 \quad (8.4)$$

$$x_1 \geq 0. \quad (8.5)$$

To convert this to standard form, first we note that

$$x_1 + x_2 \geq 3 \iff -x_1 - x_2 \leq -3 \quad (8.6)$$

which obtains the following system:

$$\min_{\vec{x} \in \mathbb{R}^2} 2x_1 + 4x_2 \quad (8.7)$$

$$\text{s.t. } -x_1 - x_2 \leq -3 \quad (8.8)$$

$$3x_1 + 2x_2 = 14 \quad (8.9)$$

$$x_1 \geq 0. \quad (8.10)$$

Adding a slack variable $x_3 \geq 0$ to the first constraint yields equality:

$$\min_{\vec{x} \in \mathbb{R}^3} 2x_1 + 4x_2 \quad (8.11)$$

$$\text{s.t. } -x_1 - x_2 + x_3 = -3 \quad (8.12)$$

$$3x_1 + 2x_2 = 14 \quad (8.13)$$

$$x_1 \geq 0 \quad (8.14)$$

$$x_3 \geq 0. \quad (8.15)$$

The only reason this is not in standard form is that x_2 is unconstrained. We can represent any real number as the difference of non-negative numbers; one such construction is $x_2 = x_2^+ - x_2^-$ where $x_2^+ = \max\{x_2, 0\}$ and $x_2^- = -\min\{x_2, 0\}$, but there are many others. Thus we can replace x_2 by $x_4 - x_5$ and add the constraints that $x_4 \geq 0$ and $x_5 \geq 0$, obtaining the problem

$$\min_{\vec{x} \in \mathbb{R}^5} 2x_1 + 4x_4 - 4x_5 \quad (8.16)$$

$$\text{s.t. } -x_1 + x_3 - x_4 + x_5 = -3 \quad (8.17)$$

$$3x_1 + 2x_4 - 2x_5 = 14 \quad (8.18)$$

$$x_1 \geq 0 \quad (8.19)$$

$$x_3 \geq 0 \quad (8.20)$$

$$x_4 \geq 0 \quad (8.21)$$

$$x_5 \geq 0. \quad (8.22)$$

Finally, we notice that x_2 is neither in the objective nor the constraints and so can be eliminated.

$$\min_{\vec{x} \in \mathbb{R}^4} 2x_1 + 4x_3 - 4x_4 \quad (8.23)$$

$$\text{s.t.} \quad -x_1 + x_2 - x_3 + x_4 = -3 \quad (8.24)$$

$$3x_1 + 2x_3 - 2x_4 = 14 \quad (8.25)$$

$$x_1 \geq 0 \quad (8.26)$$

$$x_2 \geq 0 \quad (8.27)$$

$$x_3 \geq 0 \quad (8.28)$$

$$x_4 \geq 0. \quad (8.29)$$

Linear programs are convex problems, and in particular the feasible set of a linear program is a convex set.

Proposition 183

Any linear program is a convex optimization problem.

Proposition 184

The dual of the standard form linear program

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.30)$$

$$\text{s.t.} \quad A\vec{x} = \vec{y}$$

$$\vec{x} \geq \vec{0}$$

is

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^n \\ \vec{\nu} \in \mathbb{R}^m}} -\vec{y}^\top \vec{\nu} \quad (8.31)$$

$$\text{s.t.} \quad \vec{c} - \vec{\lambda} + A^\top \vec{\nu} = \vec{0}$$

$$\vec{\lambda} \geq \vec{0}.$$

Proof. The Lagrangian is

$$L(\vec{x}, \vec{\lambda}, \vec{\nu}) = \vec{c}^\top \vec{x} - \vec{\lambda}^\top \vec{x} + \vec{\nu}^\top (A\vec{x} - \vec{y}) \quad (8.32)$$

$$= (\vec{c} - \vec{\lambda} + A^\top \vec{\nu})^\top \vec{x} - \vec{\nu}^\top \vec{y}. \quad (8.33)$$

The dual function is

$$g(\vec{\lambda}, \vec{\nu}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{\nu}) \quad (8.34)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ (\vec{c} - \vec{\lambda} + A^\top \vec{\nu})^\top \vec{x} - \vec{\nu}^\top \vec{y} \right\} \quad (8.35)$$

$$= \min_{\vec{x} \in \mathbb{R}^n} \left\{ (\vec{c} - \vec{\lambda} + A^\top \vec{\nu})^\top \vec{x} \right\} - \vec{\nu}^\top \vec{y} \quad (8.36)$$

$$= \begin{cases} -\vec{\nu}^\top \vec{y}, & \text{if } \vec{c} - \vec{\lambda} + A^\top \vec{\nu} = \vec{0} \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.37)$$

The dual problem is

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^n \\ \vec{\nu} \in \mathbb{R}^m}} g(\vec{\lambda}, \vec{\nu}) \quad (8.38)$$

$$\text{s.t. } \vec{\lambda} \geq \vec{0}. \quad (8.39)$$

Expanding this out and adding the domain constraint, we get

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^n \\ \vec{\nu} \in \mathbb{R}^m}} -\vec{\nu}^\top \vec{y} \quad (8.40)$$

$$\text{s.t. } \vec{c} - \vec{\lambda} + A^\top \vec{\nu} = \vec{0} \quad (8.41)$$

$$\vec{\lambda} \geq \vec{0} \quad (8.42)$$

as desired. \square

One very relevant question is how to solve linear programs efficiently. There are several methods available to us — for example, any constrained convex optimization solver, such as projected gradient descent, will solve the problem, given an appropriate learning rate and efficient projection method. The algorithm most used in practice is the interior point method; we will learn the basics of interior point methods later in this course.¹ But there is one algorithm which was used for many years in practice, that truly exploits the structure of linear programs to efficiently solve them. It is called the *simplex algorithm*, which was invented by George Dantzig in 1947.

The core idea behind the the simplex algorithm is that *at least one optimal point of a linear program is a “vertex” of its feasible set, so long as this feasible set is bounded*. We need to prove this idea, so far stated informally, and to do this we will first need to characterize the feasible set of a linear program.

Definition 185 (Polyhedron, Polygon)

A *polyhedron* is an intersection of a finite number of half-spaces. A *polygon* is a bounded polyhedron.

From the definition of a linear program, we see that its feasible region must be a polyhedron. For the standard form linear program, we can check explicitly that the feasible set is the intersection of the three classes of half-spaces $\{\vec{x} \in \mathbb{R}^n \mid \vec{a}_i^\top \vec{x} \leq y_i\}$, $\{\vec{x} \in \mathbb{R}^n \mid \vec{a}_i^\top \vec{x} \geq y_i\}$, and $\{x_j \in \mathbb{R}^n \mid x_j \geq 0\}$, where \vec{a}_i^\top are the rows of A , y_i are the entries of \vec{y} , and x_j are the entries of \vec{x} . This feasible region can be unbounded or bounded; if it is bounded, it will be a polygon.

Definition 186 (Extreme Point, Vertex)

Let $K \subseteq \mathbb{R}^n$ be a set. We say that $\vec{x} \in K$ is an *extreme point* of K if there *does not* exist $\vec{y}, \vec{z} \in K \setminus \{\vec{x}\}$ and $\theta \in [0, 1]$ such that $\vec{x} = \theta \vec{y} + (1 - \theta) \vec{z}$.

An extreme point of a polyhedron is called a *vertex* of the polyhedron.

If the feasible set of a linear program is an unbounded polyhedron, then there are examples where the optimal value is not achieved at a vertex, as demonstrated in the following example.

Example 187. Consider the following linear program over variables $\vec{x} \in \mathbb{R}^2$:

$$\min_{x_1, x_2 \in \mathbb{R}} \begin{bmatrix} 0 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (8.43)$$

¹Talk to your TA Tarun if you want to learn more about interior point methods!

$$\begin{aligned} \text{s.t. } x_1 &\geq 0 \\ x_2 &\geq 0 \\ x_1 &= 1. \end{aligned}$$

Define $\vec{x}_\alpha \doteq \begin{bmatrix} 1 \\ \alpha \end{bmatrix}$. One can check that \vec{x}_α is feasible for all $\alpha \geq 0$, and that the objective value at \vec{x}_α is $-\alpha$. By sending $\alpha \rightarrow \infty$, we get $p^* = -\infty$, and the optimum is not achieved at a vertex (or indeed by any $\vec{x} \in \mathbb{R}^2$).

With this understanding, we now seek to prove the main idea we had earlier. There are several proofs, but the cleanest uses the following intuitive fact:

Proposition 188

A polygon has finitely many vertices and is the convex hull of its vertices.

Unfortunately, the proof is (surprisingly?) quite complicated, so we omit it. A complete proof is in [ziegler2012lectures](#), for example.

Theorem 189 (Main Theorem of Linear Programming)

Consider the standard form linear program:

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \\ \text{s.t. } A\vec{x} &= \vec{y} \\ \vec{x} &\geq \vec{0}. \end{aligned} \tag{8.44}$$

Suppose that the feasible set $\Omega \doteq \{\vec{x} \in \mathbb{R}^n \mid A\vec{x} = \vec{y}, \vec{x} \geq \vec{0}\}$ is bounded. Then the optimal value is achieved at a vertex.

Namely, one can find an optimal point which is a vertex. There may be optimal points that are not vertices. The simplest example is to set the objective as $\vec{c} = \vec{0}$, so every feasible point is optimal with objective value 0, but there are other examples which are a bit more complicated to set up.

Proof. Since the feasible set Ω is a bounded polyhedron, it is a polygon, and so it is the convex hull of its vertices, say $\vec{v}_1, \dots, \vec{v}_k$. Thus any $\vec{x} \in \Omega$ can be written as a convex combination of the vertices \vec{v}_i , namely,

$$\vec{x} = \sum_{i=1}^k \alpha_i \vec{v}_i \tag{8.45}$$

where each $\alpha_i \geq 0$ and $\sum_{i=1}^k \alpha_i = 1$. It then follows that

$$p^* = \min_{\alpha_1, \dots, \alpha_k \in \mathbb{R}} \sum_{i=1}^k \alpha_i (\vec{c}^\top \vec{v}_i) \tag{8.46}$$

$$\text{s.t. } \alpha_i \geq 0, \quad \forall i \in \{1, \dots, k\} \tag{8.47}$$

$$\sum_{i=1}^k \alpha_i = 1. \tag{8.48}$$

Now we have

$$\sum_{i=1}^k \alpha_i (\vec{c}^\top \vec{v}_i) \geq \sum_{i=1}^k \alpha_i \left(\min_{j \in \{1, \dots, k\}} \vec{c}^\top \vec{v}_j \right) \quad (8.49)$$

$$= \left(\min_{j \in \{1, \dots, k\}} \vec{c}^\top \vec{v}_j \right) \underbrace{\left(\sum_{i=1}^k \alpha_i \right)}_{=1} \quad (8.50)$$

$$= \min_{j \in \{1, \dots, k\}} \vec{c}^\top \vec{v}_j. \quad (8.51)$$

Let $i^* \in \{1, \dots, k\}$ be an index such that $\vec{c}^\top \vec{v}_{i^*}$ achieves the above minimum, i.e., $\vec{c}^\top \vec{v}_{i^*} = \min_{j \in \{1, \dots, k\}} \vec{c}^\top \vec{v}_j$. Then the above lower bound is achieved when $\alpha_{i^*} = 1$ and $\alpha_i = 0$ for $i \neq i^*$, for example. Thus \vec{v}_{i^*} is an optimal point for the original linear program, concluding the proof. \square

This theorem says is that to solve a linear program, we only need to check the vertices of the constraint polyhedron. This reduces an optimization problem over \mathbb{R}^n to an optimization problem over the finite set of vertices. This reduction motivates a “greedy-like heuristic” solver for linear programs with bounded feasible sets Ω , which is called the *simplex method*. The simplex method is the following procedure:

- Start at a vertex \vec{v} of Ω .
- While there is a neighboring vertex \vec{w} with $\vec{c}^\top \vec{v} > \vec{c}^\top \vec{w}$, move to it, i.e., set $\vec{v} \leftarrow \vec{w}$.
- When there are no neighboring vertices with better optima, stop and return \vec{v} .

There are (rather more technical) modifications one can make to this algorithm to solve linear programs with unbounded feasible sets. But the main idea is just the same as gradient descent: iteratively search locally for another point with better objective value, and move to it.

8.2 Quadratic Programs

Definition 190 (Quadratic Program)

A *quadratic program* (QP) is an optimization problem with a quadratic objective and affine constraints. A *standard form quadratic program* is an optimization problem of the following form:

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} & (8.52) \\ \text{s.t.} \quad & A \vec{x} \leq \vec{y} \\ & C \vec{x} = \vec{z}, \end{aligned}$$

where $H \in \mathbb{S}^n$.

In the standard form, we do not lose any generality by enforcing $H \in \mathbb{S}^n$. In particular, for any H we have

$$\frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} = \frac{1}{2} \vec{x}^\top \left(\frac{H + H^\top}{2} \right) \vec{x} + \vec{c}^\top \vec{x} \quad (8.53)$$

whence the matrix $\frac{H+H^\top}{2}$ (i.e., the *symmetric part* of H) is always symmetric. So if we have a non-symmetric H we can just replace it with its symmetric part, and thus obtain a standard form quadratic program.

Quadratic programs may or may not be convex.

Proposition 191

Consider the following standard form quadratic program:

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} & (8.54) \\ \text{s.t. } & A \vec{x} \leq \vec{y} \\ & C \vec{x} = \vec{z}, \end{aligned}$$

where $H \in \mathbb{S}^n$. The following are equivalent:

- (a) The problem is convex.
- (b) $H \in \mathbb{S}_+^n$.

Proof. The gradient and Hessian of the objective are

$$\nabla \left(\frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \right) = \frac{1}{2} (H + H^\top) \vec{x} + \vec{c} = H \vec{x} + \vec{c} \quad (8.55)$$

$$\nabla^2 \left(\frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \right) = H \quad (8.56)$$

so the objective function is convex if and only if $H \in \mathbb{S}_+^n$. Thus the problem is a convex problem if and only if $H \in \mathbb{S}_+^n$ and the constraint set is convex. But the constraint set is always convex because it is defined by a set of linear equations and inequalities. Thus the problem is convex if and only if $H \in \mathbb{S}_+^n$. \square

Example 192. Let us consider the unconstrained quadratic program

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \left(\frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \right) \quad (8.57)$$

where $H \in \mathbb{S}^n$, and solve it in a variety of cases, namely that where $H \notin \mathbb{S}_+^n$, $H \in \mathbb{S}_+^n$ with $\vec{c} \in \mathcal{N}(H) \setminus \{\vec{0}\}$, and $H \in \mathbb{S}_+^n$ with $\vec{c} \in \mathcal{N}(H)^\perp = \mathcal{R}(H^\top) = \mathcal{R}(H)$.

Case 1. Suppose that $H \notin \mathbb{S}_+^n$. Then H has a negative eigenvalue λ ; let \vec{v} be any corresponding unit eigenvector.

Then, if we choose $\vec{x}_t = t \cdot \vec{v}$, we get

$$p^* \leq \lim_{t \rightarrow \infty} \left(\frac{1}{2} \vec{x}_t^\top H \vec{x}_t + \vec{c}^\top \vec{x}_t \right) \quad (8.58)$$

$$= \lim_{t \rightarrow \infty} \left(\frac{1}{2} t^2 \vec{v}^\top H \vec{v} + t \vec{c}^\top \vec{v} \right). \quad (8.59)$$

To bound the terms inside the limit, we compute

$$\vec{v}^\top H \vec{v} = \vec{v}^\top (\lambda \vec{v}) = \lambda \vec{v}^\top \vec{v} = \lambda \quad (8.60)$$

$$\vec{c}^\top \vec{v} \leq \|\vec{c}\|_2 \|\vec{v}\|_2 = \|\vec{c}\|_2. \quad (8.61)$$

Thus we have

$$p^* \leq \lim_{t \rightarrow \infty} \left(\frac{1}{2} \lambda t^2 + t \|\vec{c}\|_2 \right). \quad (8.62)$$

Since $\lambda < 0$, the term inside the limit is a concave (i.e., downward facing) quadratic function of t , and so its limit as $t \rightarrow \pm\infty$ is $-\infty$. Thus $p^* \leq -\infty$ so $p^* = -\infty$.

Case 2. Suppose that $H \in \mathbb{S}_+^n$, and suppose that $\vec{c} \in \mathcal{N}(H) \setminus \{0\}$. Then H has (at least) an eigenvalue λ equal to 0, and in particular, by the spectral theorem \vec{c} can be written as a linear combination of eigenvectors of H with eigenvalue 0. Let \vec{v} be any unit eigenvector with eigenvalue 0 such that $\vec{c}^\top \vec{v} \neq 0$. Let $\vec{x}_t = -t \cdot \text{sgn}(\vec{c}^\top \vec{v}) \cdot \vec{v}$. Then

$$p^* \leq \lim_{t \rightarrow \infty} \left(\frac{1}{2} \vec{x}_t^\top H \vec{x}_t + \vec{c}^\top \vec{x}_t \right) \quad (8.63)$$

$$= \lim_{t \rightarrow \infty} \left(\frac{1}{2} t^2 \vec{v}^\top H \vec{v} - t \cdot \text{sgn}(\vec{c}^\top \vec{v}) \cdot \vec{c}^\top \vec{v} \right) \quad (8.64)$$

$$= \lim_{t \rightarrow \infty} \left(\frac{1}{2} t^2 \vec{v}^\top \vec{0} - t \cdot |\vec{c}^\top \vec{v}| \right) \quad (8.65)$$

$$= \lim_{t \rightarrow \infty} (-t \cdot |\vec{c}^\top \vec{v}|) \quad (8.66)$$

$$= \lim_{t \rightarrow \infty} t \cdot \underbrace{(-|\vec{c}^\top \vec{v}|)}_{<0} \quad (8.67)$$

$$= -\infty. \quad (8.68)$$

Thus $p^* \leq -\infty$ so $p^* = -\infty$.

Case 3. Suppose that $H \in \mathbb{S}_+^n$ with $\vec{c} \in \mathcal{R}(H)$. Then there is nonzero \vec{x}_0 such that $\vec{c} = -H\vec{x}_0$. Rewriting the objective, we obtain

$$\frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} = \frac{1}{2} \vec{x}^\top H \vec{x} - \vec{x}_0^\top H \vec{x} \quad (8.69)$$

$$= \frac{1}{2} \vec{x}^\top H \vec{x} - \vec{x}_0^\top H \vec{x} + \frac{1}{2} \vec{x}_0^\top H \vec{x}_0 - \frac{1}{2} \vec{x}_0^\top H \vec{x}_0 \quad (8.70)$$

$$= \frac{1}{2} (\vec{x}^\top H \vec{x} - 2\vec{x}_0^\top H \vec{x} + \vec{x}_0^\top H \vec{x}_0) - \frac{1}{2} \vec{x}_0^\top H \vec{x}_0 \quad (8.71)$$

$$= \frac{1}{2} (\vec{x} - \vec{x}_0)^\top H (\vec{x} - \vec{x}_0) - \frac{1}{2} \vec{x}_0^\top H \vec{x}_0. \quad (8.72)$$

Since $H \in \mathbb{S}_+^n$, the minimizer is at any \vec{x} such that $\vec{x} - \vec{x}_0 \in \mathcal{N}(H)$. One can write this as $\vec{x} \in \vec{x}_0 + \mathcal{N}(H)$. A particular solution in terms of problem parameters is $\vec{x} = -H^\dagger \vec{c}$ where H^\dagger is the Moore-Penrose pseudoinverse of H . Recall that we discussed the Moore-Penrose pseudoinverse in more generality in homework where we derived the solution to the least-norm least-squares problem, but one can show that if $H = U\Lambda U^\top$

then $H^\dagger = U\Lambda^\dagger U^\top$ where Λ^\dagger is the diagonal matrix whose entries are $\Lambda_{ii}^\dagger = \begin{cases} 1/\Lambda_{ii}, & \text{if } \Lambda_{ii} \neq 0 \\ 0, & \text{if } \Lambda_{ii} = 0 \end{cases}$.

The previous example shows that we can solve unconstrained quadratic programs directly and read off the solutions. It turns out that one can transform any quadratic program with equality constraints into an unconstrained quadratic program. So really, this analysis encapsulates a huge class of quadratic programs.

Computing the dual of a quadratic program has a similar number of cases; it is an exercise which is left to homework.

Example 193 (Linear-Quadratic Regulator). Suppose we have a discrete-time dynamical system, of the form

$$\vec{x}_{t+1} = A\vec{x}_t + B\vec{u}_t, \quad \forall t \geq 0 \quad (8.73)$$

$$\vec{x}_0 = \vec{\xi}. \quad (8.74)$$

Then one can show that

$$\vec{x}_t = A^t \vec{\xi} + \sum_{k=0}^{t-1} A^{t-k-1} B \vec{u}_k. \quad (8.75)$$

For a fixed terminal time T , we want to reach goal state \vec{g} . Namely, we want to solve the problem

$$\min_{\substack{\vec{x}_0, \dots, \vec{x}_T \\ \vec{u}_0, \dots, \vec{u}_{T-1}}} \|\vec{x}_T - \vec{g}\|_2^2 + \sum_{k=0}^{T-1} \|\vec{u}_k\|_2^2 \quad (8.76)$$

$$\text{s.t. } \vec{x}_{t+1} = A\vec{x}_t + B\vec{u}_t, \quad \forall t \in \{0, 1, \dots, T-1\} \quad (8.77)$$

$$\vec{x}_0 = \vec{\xi}. \quad (8.78)$$

This is a quadratic program since the objective function is a quadratic function of each \vec{x}_t , and the constraints are affine equations relating the \vec{x}_t and the \vec{u}_t .

As a last note, problems with quadratic objectives and quadratic inequality constraints are called *quadratically constrained quadratic programs* (QCQPs). Like quadratic programs, QCQPs can be convex or non-convex.

8.3 Quadratically-Constrained Quadratic Programs

Definition 194 (Quadratically-Constrained Quadratic Program)

A *quadratically-constrained quadratic program* (QCQP) is an optimization problem with a quadratic objective and quadratic constraints. A *standard form quadratically constrained quadratic program* is an optimization problem of the following form:

$$\begin{aligned} p^* = \min_{\vec{x} \in \mathbb{R}^n} & \quad \frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \\ \text{s.t.} & \quad \frac{1}{2} \vec{x}^\top P_i \vec{x} + \vec{b}_i^\top \vec{x} + c_i \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & \quad \frac{1}{2} \vec{x}^\top Q_i \vec{x} + \vec{d}_i^\top \vec{x} + f_i = 0, \quad \forall i \in \{1, \dots, p\}, \end{aligned} \quad (8.79)$$

where $H, P_1, \dots, P_m, Q_1, \dots, Q_p \in \mathbb{S}^n$.

Proposition 195

Consider the following standard form quadratically-constrained quadratic program:

$$\begin{aligned} p^* = \min_{\vec{x} \in \mathbb{R}^n} & \quad \frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \\ \text{s.t.} & \quad \frac{1}{2} \vec{x}^\top P_i \vec{x} + \vec{b}_i^\top \vec{x} + c_i \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & \quad \frac{1}{2} \vec{x}^\top Q_i \vec{x} + \vec{d}_i^\top \vec{x} + f_i = 0, \quad \forall i \in \{1, \dots, p\}, \end{aligned} \quad (8.80)$$

where $H, P_1, \dots, P_m, Q_1, \dots, Q_p \in \mathbb{S}^n$. If $H, P_1, \dots, P_m \in \mathbb{S}_+^n$ and $Q_1 = \dots = Q_p = 0$, then the problem is convex.

Proof. Left as exercise. □

8.4 Second-Order Cone Programs

There is one more broad class of problems that we consider in this course, called *second-order cone programs* (SOCPs). They are among the broadest class of problems that we can efficiently solve using algorithms such as the interior point

method, which we may discuss later in the course.

Definition 196 (Second-Order Cone Program)

A *second-order cone program* is an optimization problem with a linear objective and affine and “second-order cone constraints”, i.e., constraints which say that an affine function of \vec{x} is contained in the second-order cone. A *standard form second-order cone program* is an optimization problem of the following form:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.81)$$

$$\text{s.t.} \quad \|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.82)$$

Second order cone constraints are strictly more broad than affine constraints; to encode an affine constraint $A_i \vec{x} = \vec{y}_i$ as a second-order cone constraint, pick the corresponding $\vec{b}_i = \vec{0}$ and $z_i = 0$. This makes the constraint $\|A_i \vec{x} - \vec{y}_i\|_2 \leq 0$ or equivalently $A_i \vec{x} = \vec{y}_i$.

Proposition 197

Second-order cone problems are convex optimization problems.

Proof. Each second-order cone constraint $\|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i$ can be alternatively formulated as constraining the tuple $(A_i \vec{x} - \vec{y}_i, \vec{b}_i^\top \vec{x} + z_i) \in \mathbb{R}^{n+1}$ to lie within the second-order cone in \mathbb{R}^{n+1} . But this tuple is an affine transformation of \vec{x} , in particular

$$\begin{bmatrix} A_i \vec{x} - \vec{y}_i \\ \vec{b}_i^\top \vec{x} + z_i \end{bmatrix} = \begin{bmatrix} A_i \\ \vec{b}_i^\top \end{bmatrix} \vec{x} + \begin{bmatrix} -\vec{y}_i \\ z_i \end{bmatrix}. \quad (8.83)$$

Since the second order cone is convex and the tuple is an affine transformation of \vec{x} , it follows that $\{\vec{x} \in \mathbb{R}^n \mid \|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i\}$ is a convex set. Thus the feasible set is convex (as the intersection of convex sets). The objective function is linear in \vec{x} , so the second-order cone problem is convex. \square

Example 198. Consider the following problem:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^m \|A_i \vec{x} - \vec{y}_i\|_2. \quad (8.84)$$

One can formulate this as a second-order cone program by using slack variables:

$$p^* = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{s} \in \mathbb{R}^m}} \sum_{i=1}^m s_i \quad (8.85)$$

$$\text{s.t.} \quad \|A_i \vec{x} - \vec{y}_i\|_2 \leq s_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.86)$$

The following problem is also very related:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \max_{i \in \{1, \dots, m\}} \|A_i \vec{x} - \vec{y}_i\|_2. \quad (8.87)$$

We can use a similar slack variable reformulation to formulate this problem as a second order cone program.

$$p^* = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ s \in \mathbb{R}}} s \quad (8.88)$$

$$\text{s.t.} \quad \|A_i \vec{x} - \vec{y}_i\|_2 \leq s, \quad \forall i \in \{1, \dots, m\}. \quad (8.89)$$

These problems can be formulated in terms of route planning – more specifically, finding the route which minimizes the total length between waypoints (in the first problem), or the route which minimizes the maximum length between waypoints (in the second problem).

Example 199 (LPs, QPs, and QCQPs as SOCPs). One can see how LPs are QPs and how QPs are QCQPs, because in each transition the set of properties becomes more permissive — first a linear objective can become a quadratic objective, then linear constraints can become quadratic constraints. It is less clear how LPs, QPs, and QCQPs are SOCPs. In this example we derive a way to write QCQPs as SOCPs, which is also applicable to LPs and QPs (since LPs and QPs are QCQPs).

Consider a QCQP of the form

$$\begin{aligned} p^* &= \min_{\vec{x} \in \mathbb{R}^n} \frac{1}{2} \vec{x}^\top H \vec{x} + \vec{c}^\top \vec{x} \\ \text{s.t.} \quad & \frac{1}{2} \vec{x}^\top P_i \vec{x} + \vec{b}_i^\top \vec{x} + c_i \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & C \vec{x} = \vec{z}, \end{aligned} \tag{8.90}$$

where $H, P_1, \dots, P_m \in \mathbb{S}_+^n$. We use the epigraph reformulation to obtain

$$\begin{aligned} p^* &= \min_{\substack{\vec{x} \in \mathbb{R}^n \\ t \in \mathbb{R}}} t + \vec{c}^\top \vec{x} \\ \text{s.t.} \quad & \frac{1}{2} \vec{x}^\top P_i \vec{x} + \vec{b}_i^\top \vec{x} + c_i \leq 0, \quad \forall i \in \{1, \dots, m\} \\ & \frac{1}{2} \vec{x}^\top H \vec{x} \leq t \\ & C \vec{x} = \vec{z}. \end{aligned} \tag{8.91}$$

Let us try to convert the quadratic constraint

$$\frac{1}{2} \vec{x}^\top P_i \vec{x} + \vec{b}_i^\top \vec{x} + c_i \leq 0 \tag{8.92}$$

into a second-order cone constraint. Notice that this constraint is equivalent to

$$\vec{x}^\top P_i \vec{x} + 2(\vec{b}_i^\top \vec{x} + c_i) \leq 0. \tag{8.93}$$

In order to write each term as a square, we first write $\vec{x}^\top P_i \vec{x} = \left\| P_i^{1/2} \vec{x} \right\|_2^2$. We now use a difference of squares identity:

$$(u + v)^2 - (u - v)^2 = 4uv. \tag{8.94}$$

To apply this to the above formula, we plug in $u = \frac{1}{2}$ and $v = \vec{b}_i^\top \vec{x} + c_i$, to get

$$2(\vec{b}_i^\top \vec{x} + c_i) = \left(\frac{1}{2} + \vec{b}_i^\top \vec{x} + c_i \right)^2 - \left(\frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i \right)^2. \tag{8.95}$$

This gives us the constraint

$$\left\| P_i^{1/2} \vec{x} \right\|_2^2 + \left(\frac{1}{2} + \vec{b}_i^\top \vec{x} + c_i \right)^2 - \left(\frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i \right)^2 \leq 0, \tag{8.96}$$

which is equivalent to

$$\left\| P_i^{1/2} \vec{x} \right\|_2^2 + \left(\frac{1}{2} + \vec{b}_i^\top \vec{x} + c_i \right)^2 \leq \left(\frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i \right)^2. \tag{8.97}$$

Now we would like to take square roots and write things in terms of the ℓ^2 -norm. For this, we need to show that $\frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i \geq 0$. This follows because, since H is positive semidefinite, we have $\frac{1}{2} \geq 0 \geq -\frac{1}{2} \vec{x}^\top H \vec{x}$, and so

$$\frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i \geq -\frac{1}{2} \vec{x}^\top H \vec{x} - \vec{b}_i^\top \vec{x} - c_i \geq 0. \quad (8.98)$$

Now, taking square roots and writing things in terms of the ℓ^2 -norm, we have

$$\left\| \begin{bmatrix} P_i^{1/2} \vec{x} \\ \frac{1}{2} + \vec{b}_i^\top \vec{x} + c_i \end{bmatrix} \right\|_2 \leq \frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i. \quad (8.99)$$

Thus we can write the QCQP as

$$\begin{aligned} p^* &= \min_{\substack{\vec{x} \in \mathbb{R}^n \\ t \in \mathbb{R}}} t + \vec{c}^\top \vec{x} \\ \text{s.t.} \quad & \left\| \begin{bmatrix} P_i^{1/2} \vec{x} \\ \frac{1}{2} + \vec{b}_i^\top \vec{x} + c_i \end{bmatrix} \right\|_2 \leq \frac{1}{2} - \vec{b}_i^\top \vec{x} - c_i, \quad \forall i \in \{1, \dots, m\} \\ & \left\| \begin{bmatrix} H^{1/2} \vec{x} \\ \frac{1}{2} - t \end{bmatrix} \right\|_2 \leq \frac{1}{2} + t \\ & \|C\vec{x} - \vec{z}\|_2 \leq 0. \end{aligned} \quad (8.100)$$

Below, we establish that the dual of an SOCP is an SOCP. This fact can either be proved via conic duality, or proved directly.

Theorem 200

Let $\vec{c} \in \mathbb{R}^n$, and for $i \in \{1, \dots, m\}$ let $A_i \in \mathbb{R}^{d_i \times n}$, $\vec{y}_i \in \mathbb{R}^{d_i}$, $\vec{b}_i \in \mathbb{R}^n$, and $z_i \in \mathbb{R}$. The dual of the following SOCP in standard form:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.101)$$

$$\text{s.t.} \quad \|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.102)$$

can be formulated as an SOCP in standard form.

Proof via Conic Duality. Let $K_i \doteq \{(\vec{u}, r) \in \mathbb{R}^{d_i} \times \mathbb{R} \mid \|\vec{u}\|_2 \leq r\}$ denote the second-order cone in \mathbb{R}^{d_i+1} , and let $d = \sum_{i=1}^m d_i$. Then the standard form SOCP can be written as:

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.103)$$

$$\text{s.t.} \quad -(A_i \vec{x} - \vec{y}_i, \vec{b}_i^\top \vec{x} + z_i) \preceq_{K_i} \vec{0}, \quad \forall i \in \{1, \dots, m\}. \quad (8.104)$$

The Lagrangian $L: \mathbb{R}^n \times \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$ can thus be defined as follows. For each $\vec{x} \in \mathbb{R}^n$, $\vec{\lambda} = (\vec{\lambda}_1, \dots, \vec{\lambda}_m) \in \mathbb{R}^d$ (with $\vec{\lambda}_i \in \mathbb{R}^{d_i}$ for each $i \in \{1, \dots, m\}$), and $\vec{\mu} \in \mathbb{R}^m$:

$$L(\vec{x}, \vec{\lambda}, \vec{\mu}) = \vec{c}^\top \vec{x} - \sum_{i=1}^m \left[\vec{\lambda}_i^\top (A_i \vec{x} - \vec{y}_i) + \mu_i (\vec{b}_i^\top \vec{x} + z_i) \right] \quad (8.105)$$

$$= \left(\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) \right)^\top \vec{x} + \sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i). \quad (8.106)$$

Next, define the dual function $g: \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$ by maximizing over the primal variable $\vec{x} \in \mathbb{R}^n$:

$$g(\vec{\lambda}, \vec{\mu}) = \min_{\vec{x} \in \mathbb{R}^n} L(\vec{x}, \vec{\lambda}, \vec{\mu}) \quad (8.107)$$

$$= \begin{cases} \sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i), & \text{if } \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) = \vec{c}, \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.108)$$

The last equality follows by noticing that the objective is linear in \vec{x} ; if its coefficient $\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) = \vec{0}$, then the objective value is the sum $\sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i)$ regardless of the value of \vec{x} , while if $\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) \neq \vec{0}$ then we can make the objective value as low as we want by picking \vec{x} appropriately. For instance, let $K > 0$ be a large positive number; then

$$\vec{x} = -K \left(\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) \right) \implies L(\vec{x}, \vec{\lambda}, \vec{\mu}) = -K \left\| \vec{c} - \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) \right\|_2^2 + \sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i) \quad (8.109)$$

which we can drive down to $-\infty$ by increasing K to $+\infty$. Thus, the dual problem is given by:

$$\max_{\substack{\vec{\lambda} \in \mathbb{R}^d \\ \vec{\mu} \in \mathbb{R}^m}} \sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i) \quad (8.110)$$

$$\text{s.t. } \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) = \vec{c}, \quad (8.111)$$

$$- (\vec{\lambda}_i, \mu_i) \preceq_{K_i} \mathbf{0}, \quad \forall i \in \{1, \dots, m\}. \quad (8.112)$$

where the last line uses the fact that for each K_i its conic dual is itself, or in standard SOCP form by:

$$\max_{\substack{\vec{\lambda} \in \mathbb{R}^d \\ \vec{\mu} \in \mathbb{R}^m}} \sum_{i=1}^m (\vec{\lambda}_i^\top \vec{y}_i - \mu_i z_i) \quad (8.113)$$

$$\text{s.t. } \left\| \sum_{i=1}^m (A_i^\top \vec{\lambda}_i + \mu_i \vec{b}_i) - \vec{c} \right\|_2 \leq 0 \quad (8.114)$$

$$\left\| \vec{\lambda}_i \right\|_2 \leq \mu_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.115)$$

□

Direct Proof. We re-iterate the SOCP to take the dual of:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.116)$$

$$\text{s.t. } \|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.117)$$

We add some variables to simplify. Namely, we introduce $\vec{u}_i \in \mathbb{R}^{d_i}$ and $w_i \in \mathbb{R}$ for each $i \in \{1, \dots, m\}$. For convenience, we define $\vec{u} \doteq (\vec{u}_1, \dots, \vec{u}_m) \in \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_m} = \mathbb{R}^d$, where again $d = \sum_{i=1}^m d_i$, and also define $\vec{w} = (w_1, \dots, w_m) \in \mathbb{R}^m$. With these definitions, the SOCP can be written as

$$p^* = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{u} \in \mathbb{R}^d \\ \vec{w} \in \mathbb{R}^m}} \vec{c}^\top \vec{x} \quad (8.118)$$

$$\text{s.t. } \|\vec{u}_i\|_2 \leq w_i, \quad \forall i \in \{1, \dots, m\} \quad (8.119)$$

$$\vec{u}_i = A_i \vec{x} - \vec{y}_i, \quad \forall i \in \{1, \dots, m\} \quad (8.120)$$

$$w_i = \vec{b}_i^\top \vec{x} + z_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.121)$$

We can thus define a Lagrangian for this system, say with dual variables $\vec{\lambda} \in \mathbb{R}^m$, $\vec{\eta} \in \mathbb{R}^d$ (with $\vec{\eta}_i \in \mathbb{R}^{d_i}$ for each i) and $\vec{\nu} \in \mathbb{R}^m$. We have

$$L(\vec{x}, \vec{u}, \vec{w}, \vec{\lambda}, \vec{\eta}, \vec{\nu}) = \vec{c}^\top \vec{x} + \sum_{i=1}^m \lambda_i (\|\vec{u}_i\|_2 - w_i) + \sum_{i=1}^m \vec{\eta}_i^\top (\vec{u}_i - A_i \vec{x} + \vec{y}_i) + \sum_{i=1}^m \nu_i (w_i - \vec{b}_i^\top \vec{x} - z_i) \quad (8.122)$$

$$= \left(\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \right)^\top \vec{x} + \sum_{i=1}^m (\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i) + \sum_{i=1}^m (-\lambda_i + \nu_i) w_i + \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i). \quad (8.123)$$

Now define the dual function $g: \mathbb{R}^m \times \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$ by minimizing over the primal variables $(\vec{x}, \vec{u}, \vec{w}) \in \mathbb{R}^n \times \mathbb{R}^d \times \mathbb{R}^m$:

$$g(\vec{\lambda}, \vec{\eta}, \vec{\nu}) = \min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{u} \in \mathbb{R}^d \\ \vec{w} \in \mathbb{R}^m}} L(\vec{x}, \vec{u}, \vec{w}, \vec{\lambda}, \vec{\eta}, \vec{\nu}) \quad (8.124)$$

$$= \begin{cases} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i), & \text{if } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \\ & \text{and } \|\vec{\eta}\|_2 \leq \lambda_i \quad \forall i \in \{1, \dots, m\} \\ & \text{and } \lambda_i = \nu_i \quad \forall i \in \{1, \dots, m\} \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.125)$$

The last equality looks complicated and a bit magical but we methodically justify it here.

(a) The Lagrangian is linear in \vec{x} , indeed having the form

$$L = \left(\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \right)^\top \vec{x} + \text{other terms not involving } \vec{x}, \quad (8.126)$$

so unless the coefficient $\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i)$ is $\vec{0}$, then we can make the Lagrangian arbitrarily negative by varying \vec{x} while keeping \vec{u} and \vec{w} fixed. For instance, let $K > 0$ be a large positive number. Then

$$\vec{x} = -K \left(\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \right) \implies L = -K \left\| \vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \right\|_2^2 + \text{other terms not involving } \vec{x} \quad (8.127)$$

which we can drive down to $-\infty$ by sending $K \rightarrow \infty$. On the other hand, if $\vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) = \vec{0}$, then the first term in the Lagrangian is $\vec{0}$ regardless of the value of \vec{x} .

Thus, we have proved

$$\min_{\vec{x} \in \mathbb{R}^n} L = \begin{cases} \sum_{i=1}^m (\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i) + \sum_{i=1}^m (-\lambda_i + \nu_i) w_i + \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i), & \text{if } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.128)$$

(b) Suppose that $\vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i)$. The Lagrangian has the form

$$\min_{\vec{x} \in \mathbb{R}^n} L = \sum_{i=1}^m (\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i) + \text{other terms not involving } \vec{u}. \quad (8.129)$$

Towards minimizing this expression over \vec{u} , we aim to solve the problem

$$\min_{\vec{u}_i \in \mathbb{R}^{d_i}} (\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i), \quad (8.130)$$

and collect the results for each i at the end. At first glance, it may seem hard to imagine this term blowing up at all. Towards finding out a possible blow-up case, if any, we use Cauchy-Schwarz to try to make the sum as small as possible. In particular, by Cauchy-Schwarz we have

$$\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i \geq \lambda_i \|\vec{u}_i\|_2 - \|\vec{\eta}_i\|_2 \|\vec{u}_i\|_2 = (\lambda_i - \|\vec{\eta}_i\|_2) \|\vec{u}_i\|_2 \quad (8.131)$$

with equality when \vec{u}_i points in the opposite direction as $\vec{\eta}_i$, that is, $\vec{u}_i = -K\vec{\eta}_i$ for some $K \geq 0$. With this value of \vec{u}_i (for varying $K \rightarrow \infty$) we shall try to make the Lagrangian go to $-\infty$. Indeed, in this case, we have

$$\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i = K(\lambda_i - \|\vec{\eta}_i\|_2) \|\vec{\eta}_i\|_2. \quad (8.132)$$

First suppose that $\|\vec{\eta}_i\|_2 = 0$. Since $\lambda_i \geq 0$ in the Lagrangian formulation, we must have $\lambda_i \geq \|\vec{\eta}_i\|_2$, as indicated in the original equality. The optimal \vec{u}_i is $\vec{u}_i = -K\vec{\eta}_i = \vec{0}$ (independently of the value of K), at which point the term in the Lagrangian becomes 0. We now deal with the non-edge case, assuming that $\vec{\eta}_i \neq \vec{0}$.

Suppose that $\lambda_i - \|\vec{\eta}_i\|_2 < 0$. Then by sending $K \rightarrow \infty$ with this choice of $\vec{u}_i = -K\vec{\eta}_i$ we drive the Lagrangian to $-\infty$. On the other hand, if $\lambda_i - \|\vec{\eta}_i\|_2 \geq 0$, then the minimizing choice for K is $K = 0$, so that $\vec{u}_i = \vec{0}$, and the term in the Lagrangian becomes 0. Thus,

$$\min_{\vec{u}_i \in \mathbb{R}^{d_i}} (\lambda_i \|\vec{u}_i\|_2 + \vec{\eta}_i^\top \vec{u}_i) = \begin{cases} 0, & \text{if } \lambda_i \geq \|\vec{\eta}_i\|_2 \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.133)$$

Applying this logic to each $i \in \{1, \dots, m\}$, we obtain

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{u} \in \mathbb{R}^d}} L = \begin{cases} \sum_{i=1}^m (-\lambda_i + \nu_i) w_i + \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i), & \text{if } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \\ & \text{and } \lambda_i \geq \|\vec{\eta}_i\|_2, \quad \forall i \in \{1, \dots, m\} \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.134)$$

(c) Suppose that $\vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i)$ and $\lambda_i \geq \|\vec{\eta}_i\|_2$ for each i . Then the Lagrangian has the form

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{u} \in \mathbb{R}^d}} L = \sum_{i=1}^m (\nu_i - \lambda_i) w_i + \text{other terms not involving } \vec{w}. \quad (8.135)$$

Towards minimizing this expression over \vec{w} , we aim to solve the problem

$$\min_{w_i \in \mathbb{R}} (\nu_i - \lambda_i) w_i \quad (8.136)$$

and collect the results at the end. Thankfully this is much simpler than the rest of the calculations, since the objective is an unconstrained minimization of a linear function of a scalar w_i . If the coefficient $\nu_i - \lambda_i$ is nonzero, then we can thus blow up the objective in any direction by choosing w_i accordingly. Namely, if $\nu_i \neq \lambda_i$ then the choice of $w_i = -K(\nu_i - \lambda_i)$ for some positive scalar $K > 0$, simplifies the objective as $-K(\nu_i - \lambda_i)^2$. Since $(\nu_i - \lambda_i)^2 > 0$, taking $K \rightarrow \infty$ shows that the optimal value of the objective is $-\infty$. On the other hand, if $\nu_i = \lambda_i$ then the objective has value 0 independent of the choice of w_i . We have shown that

$$\min_{w_i \in \mathbb{R}} (\nu_i - \lambda_i) w_i = \begin{cases} 0, & \text{if } \nu_i = \lambda_i \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.137)$$

Applying this logic to all $i \in \{1, \dots, m\}$, we obtain

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{u} \in \mathbb{R}^d \\ \vec{w} \in \mathbb{R}^m}} L = \begin{cases} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i), & \text{if } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \\ & \text{and } \lambda_i \geq \|\vec{u}_i\|_2, \quad \forall i \in \{1, \dots, m\} \\ & \text{and } \nu_i = \lambda_i, \quad \forall i \in \{1, \dots, m\} \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.138)$$

Now we can write down the dual problem as

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}_+^m \\ \vec{\eta} \in \mathbb{R}^d \\ \vec{\nu} \in \mathbb{R}^m}} g(\vec{\lambda}, \vec{\eta}, \vec{\nu}) \quad (8.139)$$

which simplifies to

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^m \\ \vec{\eta} \in \mathbb{R}^d \\ \vec{\nu} \in \mathbb{R}^m}} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i) \quad (8.140)$$

$$\text{s.t. } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \quad (8.141)$$

$$\lambda_i \geq \|\vec{u}_i\|_2, \quad \forall i \in \{1, \dots, m\} \quad (8.142)$$

$$\lambda_i = \nu_i, \quad \forall i \in \{1, \dots, m\} \quad (8.143)$$

$$\lambda_i \geq 0, \quad \forall i \in \{1, \dots, m\}. \quad (8.144)$$

Note that the constraint $\lambda_i \geq \|\vec{u}_i\|_2$ already implies $\lambda_i \geq 0$ since $\|\vec{u}_i\|_2 \geq 0$. Thus, we can rewrite the problem again as

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^m \\ \vec{\eta} \in \mathbb{R}^d \\ \vec{\nu} \in \mathbb{R}^m}} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \nu_i z_i) \quad (8.145)$$

$$\text{s.t. } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \nu_i \vec{b}_i) \quad (8.146)$$

$$\lambda_i \geq \|\vec{u}_i\|_2, \quad \forall i \in \{1, \dots, m\} \quad (8.147)$$

$$\lambda_i = \nu_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.148)$$

Now note that the last constraint forces $\vec{\lambda} = \vec{\nu}$. Thus, we can eliminate one of them; we choose arbitrarily to eliminate $\vec{\nu}$ by replacing it everywhere with $\vec{\lambda}$. This gives the dual problem as

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^m \\ \vec{\eta} \in \mathbb{R}^d}} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \lambda_i z_i) \quad (8.149)$$

$$\text{s.t. } \vec{c} = \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \lambda_i \vec{b}_i) \quad (8.150)$$

$$\lambda_i \geq \|\vec{u}_i\|_2, \quad \forall i \in \{1, \dots, m\}. \quad (8.151)$$

To write this in SOCP form, we can write the affine constraint as a norm, obtaining

$$d^* = \max_{\substack{\vec{\lambda} \in \mathbb{R}^m \\ \vec{\eta} \in \mathbb{R}^d}} \sum_{i=1}^m (\vec{\eta}_i^\top \vec{y}_i - \lambda_i z_i) \quad (8.152)$$

$$\text{s.t.} \quad \left\| \vec{c} - \sum_{i=1}^m (A_i^\top \vec{\eta}_i + \lambda_i \vec{b}_i) \right\|_2 \leq 0 \quad (8.153)$$

$$\|\vec{u}_i\|_2 \leq \lambda_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.154)$$

Thus, we have obtained that the dual of an SOCP is another SOCP. \square

The following content is optional/out of scope for this semester. Regardless, it may be helpful to read it to gain context, or get a deeper understanding of various results.

8.5 Semidefinite Programming

This section introduces the semidefinite program (SDP), one of the broadest classes of named optimization problems. We begin with its two forms, the inequality form and standard form, and their properties.

In this section, we will use \succeq and \preceq to denote inequalities between symmetric matrices. These are instances of generalized inequalities, as in Definition 176, associated with the (proper) cone of positive semidefinite matrices. All you need to know is: if we write $A \succeq 0$, this means A is symmetric positive semidefinite, whereas if $A \succeq B$ then it means $A - B$ is symmetric positive semidefinite. On the other hand, if we write $A \preceq 0$, this means $-A$ is symmetric positive semidefinite; that is, A is *symmetric negative semidefinite*, with all non-positive eigenvalues.

Recall that \mathbb{S}^n is the set of $n \times n$ symmetric matrices, and \mathbb{S}_+^n is the set of $n \times n$ symmetric positive semidefinite matrices.

Definition 201 (Semidefinite Program in Inequality Form)

A *semidefinite program in inequality form* is an optimization problem of the following form:

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.155)$$

$$\text{s.t.} \quad F_0 + \sum_{i=1}^n x_i F_i \preceq 0, \quad (8.156)$$

where $\vec{c} \in \mathbb{R}^n$, and $F_0, F_1, \dots, F_n \in \mathbb{S}^n$.

The expression $F_0 + \sum_{i=1}^n x_i F_i$ is referred to as a *linear matrix inequality*. The constraint set, i.e., the set of $\vec{x} \in \mathbb{R}^n$ such that $F_0 + \sum_{i=1}^n x_i F_i \preceq 0$, is called a *spectrahedron*.

Notice that we only require one linear matrix inequality in the definition. What if we had multiple? Suppose that we actually wanted to solve the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.157)$$

$$\text{s.t.} \quad F_0^{(1)} + \sum_{i=1}^n x_i F_i^{(1)} \preceq 0, \quad (8.158)$$

$$F_0^{(2)} + \sum_{i=1}^n x_i F_i^{(2)} \preceq 0, \quad (8.159)$$

$$\vdots \quad (8.160)$$

$$F_0^{(k)} + \sum_{i=1}^n x_i F_i^{(k)} \preceq 0. \quad (8.161)$$

This could be phrased using a single linear matrix inequality, and the problem would be

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.162)$$

$$\text{s.t.} \quad \begin{bmatrix} F_0^{(1)} & & \\ & \ddots & \\ & & F_0^{(k)} \end{bmatrix} + \sum_{i=1}^n x_i \begin{bmatrix} F_i^{(1)} & & \\ & \ddots & \\ & & F_i^{(k)} \end{bmatrix} \preceq 0. \quad (8.163)$$

(If this reduction isn't clear to you, it's totally fine; try to prove it as an exercise.)

We now introduce another major standard form of SDPs.

Definition 202 (Semidefinite Program in Standard Form)

A *semidefinite program in standard form* is an optimization problem of the following form:

$$\min_{X \in \mathbb{S}^n} \text{tr}(CX) \quad (8.164)$$

$$\text{s.t.} \quad \text{tr}(A_k X) = b_k, \quad \forall k \in \{1, \dots, m\} \quad (8.165)$$

$$X \succeq 0, \quad (8.166)$$

where $C, A_1, \dots, A_m \in \mathbb{S}^n$, and $b_1, \dots, b_m \in \mathbb{R}$.

The first main theorem below establishes that the inequality and standard forms of an SDP are equivalent, in the sense that either can be reformulated as the other.

Theorem 203

An SDP in inequality form can be reformulated as an SDP in standard form, and vice versa.

Proof. Just for this proof, we introduce the notation $\text{vec}: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{mn}$, which takes an $m \times n$ matrix and unrolls it into an mn -length vector. With this notation, in fact, for two symmetric matrices $A, B \in \mathbb{S}^n$, we can write $\text{tr}(AB) = \sum_{i=1}^n \sum_{j=1}^n A_{ij} B_{ij} = \text{vec}(A)^\top \text{vec}(B)$. On the other hand, we will sometimes need to access the element of $\text{vec}(A)$ corresponding to A_{ij} ; we denote this by $\text{vec}(A)_{i,j}$ (where the comma makes it clear that the index is not the product of i and j). We will also use the notation $\text{diag}: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ which takes a vector and returns a diagonal matrix whose diagonal is the entries of this vector. This notation will greatly simplify things to follow.

“Inequality form \implies Standard form”: Let $\vec{c} \in \mathbb{R}^n$ and $F_0, F_1, \dots, F_n \in \mathbb{S}^d$, and consider the following SDP in inequality form:

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.167)$$

$$\text{s.t.} \quad F_0 + \sum_{i=1}^n x_i F_i \preceq 0. \quad (8.168)$$

Our goal is to write it in the form

$$\min_{X \in \mathbb{S}^m} \text{tr}(CX) \quad (8.169)$$

$$\text{s.t.} \quad \text{tr}(A_k X) = b_k, \quad \forall k \in \{1, \dots, p\}, \quad (8.170)$$

$$X \succeq 0. \quad (8.171)$$

First, towards introducing the positive semidefinite constraint, we introduce a new variable $Y \in \mathbb{S}^d$, associated with $-(F_0 + \sum_{i=1}^n x_i F_i)$. That is, our original problem has the form

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ Y \in \mathbb{S}^d}} \vec{c}^\top \vec{x} \quad (8.172)$$

$$\text{s.t. } Y + F_0 + \sum_{i=1}^n x_i F_i = 0, \quad (8.173)$$

$$Y \succeq 0. \quad (8.174)$$

Since we have a linear matrix equality, we can write it as a bunch of scalar equations to get it closer to the desired form, say in the following way:

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ Y \in \mathbb{S}^d}} \vec{c}^\top \vec{x} \quad (8.175)$$

$$\text{s.t. } Y_{jk} + (F_0)_{jk} + \sum_{i=1}^n x_i (F_i)_{jk} = 0, \quad \forall j, k \in \{1, \dots, d\} \quad (8.176)$$

$$Y \succeq 0. \quad (8.177)$$

But even this isn't quite right – after all, we require all decision variables to be encapsulated in a positive semidefinite matrix. The simplest way to do this is to form a block diagonal matrix where each block is an embedding of a decision variable into a positive semidefinite matrix; the large matrix will also be positive semidefinite in this case. Towards converting \vec{x} to a positive semidefinite block, one could consider its diagonal matrix equivalent $\text{diag}(\vec{x})$, but this would not be positive semidefinite unless all entries of \vec{x} were positive. To ensure that this happens, we use slack variables, akin to the proof that general linear programs can be written in standard form.

Namely, associate vectors $\vec{x}^+, \vec{x}^- \in \mathbb{R}^n$ by the following formulae:

$$x_i^+ = \begin{cases} x_i, & x_i > 0 \\ 0, & x_i \leq 0, \end{cases} \quad x_i^- = \begin{cases} 0, & x_i > 0 \\ -x_i, & x_i < 0. \end{cases} \quad (8.178)$$

In this case $\vec{x} = \vec{x}^+ - \vec{x}^-$. Thus the original problem is equivalent to the reformulation

$$\min_{\substack{\vec{x}^+ \in \mathbb{R}^n \\ \vec{x}^- \in \mathbb{R}^n \\ Y \in \mathbb{S}^n}} \vec{c}^\top (\vec{x}^+ - \vec{x}^-) \quad (8.179)$$

$$\text{s.t. } Y_{jk} + (F_0)_{jk} + \sum_{i=1}^n (x_i^+ - x_i^-) (F_i)_{jk} = 0, \quad \forall j, k \in \{1, \dots, d\} \quad (8.180)$$

$$x_i^+ \geq 0, \quad \forall i \in \{1, \dots, n\}, \quad (8.181)$$

$$x_i^- \geq 0, \quad \forall i \in \{1, \dots, n\}, \quad (8.182)$$

$$Y \succeq 0. \quad (8.183)$$

Now we are in business; we can write all the inequality/definiteness constraints as

$$Z \doteq \begin{bmatrix} \text{diag}(\vec{x}^+) & & \\ & \text{diag}(\vec{x}^-) & \\ & & Y \end{bmatrix} \succeq 0. \quad (8.184)$$

This is the positive semidefiniteness constraint we want, so the decision variable is $Z \in \mathbb{S}^{2n+d}$. As notation, let $Z^{1,i} = \text{diag}(\vec{x}^+)_{ii} = x_i^+$ be the i^{th} element of the first block, $Z^{2,i} = \text{diag}(\vec{x}^-)_{ii} = x_i^-$ be the i^{th} element of the second block, and $Z^{3,i,j} = Y_{ij}$ be the $(i,j)^{\text{th}}$ element of the third block. As notation for later, let \mathcal{O} be the set of all indices in $\{1, \dots, 2n+d\} \times \{1, \dots, 2n+d\}$ which are not on the diagonal or part of the Y block, and thus must be set to zero; formally $\mathcal{O} = \{(i,j) \mid 1 \leq i,j \leq 2n+d, i \neq j, i \leq 2n \text{ or } j \leq 2n\}$.

Now, we have written our problem in the form

$$\min_{Z \in \mathbb{S}^{2n+d}} \sum_{i=1}^n c_i (Z^{1,i} - Z^{2,i}) \quad (8.185)$$

$$\text{s.t. } Z^{3,jk} + (F_0)_{jk} + \sum_{i=1}^n (Z^{1,i} - Z^{2,i})(F_i)_{jk} = 0, \quad \forall j, k \in \{1, \dots, d\} \quad (8.186)$$

$$Z_{i,j} = 0, \quad \forall (i,j) \in \mathcal{O}, \quad (8.187)$$

$$Z \succeq 0. \quad (8.188)$$

Notice that all constraints are affine or positive definite, and our objective is affine; by our discussion of affine functions, the affine constraints can be written in the form $\text{tr}(A_k Z) = b_k$, and the objective can be written in the form $\text{tr}(CZ)$, for some symmetric matrices A_k, C and scalars b_k , and $k \in \{1, \dots, m\}$ where $m = d^2 + |\mathcal{O}|$.² Thus we can write our problem as

$$\min_{Z \in \mathbb{S}^{2n+d}} \text{tr}(CZ) \quad (8.189)$$

$$\text{s.t. } \text{tr}(A_k Z) = b_k, \quad \forall k \in \{1, \dots, m\} \quad (8.190)$$

$$Z \succeq 0, \quad (8.191)$$

as desired.

“Standard form \implies Inequality form”: Let $C, A_1, \dots, A_m \in \mathbb{S}^n$ be fixed symmetric matrices, and let $b_1, \dots, b_m \in \mathbb{R}$ be fixed scalars. Consider the following SDP in standard form:

$$\min_{X \in \mathbb{S}^n} \text{tr}(CX) \quad (8.192)$$

$$\text{s.t. } \text{tr}(A_k X) = b_k, \quad \forall k \in \{1, \dots, m\} \quad (8.193)$$

$$X \succeq 0. \quad (8.194)$$

We want to write it in the form

$$\min_{\vec{x} \in \mathbb{R}^m} \vec{c}^\top \vec{x} \quad (8.195)$$

$$\text{s.t. } F_0 + \sum_{i=1}^m x_i F_i \preceq 0. \quad (8.196)$$

Notice by our notation that $\text{tr}(CX) = \text{vec}(C)^\top \text{vec}(X)$ and that $\text{tr}(A_k X) = \text{vec}(A_k)^\top \text{vec}(X)$. Thus, letting $\vec{x} \in \mathbb{R}^{n^2}$ be defined as $\vec{x} = \text{vec}(X)$, our objective is linear in \vec{x} , since it is $\vec{c}^\top \vec{x}$ where $\vec{c} = \text{vec}(C)$. Furthermore, our

²Careful readers may notice that the discussion on affine functions ensured something slightly different; namely, for an affine function f on symmetric matrices (or indeed all of $\mathbb{R}^{n \times n}$), there was some matrix A and scalar b such that $f(X) = \text{tr}(A^\top X) + b$. In particular, the result did not guarantee that such an A could be symmetric. But certainly the matrix $(A + A^\top)/2$ is symmetric, and for Z symmetric, we have $\text{tr}(A^\top Z) = \text{tr}([(A + A^\top)/2]Z)$, so indeed, for an affine function $f: \mathbb{S}^n \rightarrow \mathbb{R}$ there exists some matrix $A \in \mathbb{S}^n$ and scalar $b \in \mathbb{R}$ such that $f(X) = \text{tr}(AX) + b$.

equality constraints are affine, since they are $\vec{a}_k^\top \vec{x} = b_k$ where $\vec{a}_k = \text{vec}(A_k)$. We express each equality constraint as a pair of linear matrix inequalities, since that is the only type of constraint we are permitted to have. Indeed, we have

$$\vec{a}_k^\top \vec{x} = b_k \iff \sum_{i=1}^m x_i (\vec{a}_k)_i = b_k \quad (8.197)$$

$$\iff -b_k + \sum_{i=1}^m x_i (\vec{a}_k)_i = 0 \quad (8.198)$$

$$\iff -b_k + \sum_{i=1}^m x_i (\vec{a}_k)_i \preceq 0 \text{ and } -\left(-b_k + \sum_{i=1}^m x_i (\vec{a}_k)_i\right) \preceq 0, \quad (8.199)$$

where we are using \preceq for ordering on the space of 1×1 symmetric matrices, i.e., scalars. These are bona-fide linear matrix inequalities and will be combined with others, later, to form the full linear matrix inequality constraint for our problem.

The only constraint remaining that cannot easily be expressed in vectorized form is the constraint $X \succeq 0$. For this, we note that we are allowed to have a linear matrix inequality constraint, so we want to express $X \succeq 0$ in terms of a linear matrix inequality involving \vec{x} . This is difficult at first, so we handle it in the case $n = 2$ for an example. Write

$$X = \begin{bmatrix} x_1 & x_2 \\ x_3 & x_4 \end{bmatrix}, \quad \vec{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}. \quad (8.200)$$

Notice that, since X is symmetric (and so $x_2 = x_3$), we can write X in terms of a linear combination of constant symmetric matrices, as follows

$$X = \begin{bmatrix} x_1 & x_2 \\ x_2 & x_4 \end{bmatrix} \quad (8.201)$$

$$= x_1 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + x_2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + x_4 \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad (8.202)$$

$$= x_1 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + \frac{1}{2}x_2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + \frac{1}{2}x_3 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + x_4 \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad (8.203)$$

$$= x_1 E^{11} + \frac{1}{2}x_2 (E^{12} + E^{21}) + \frac{1}{2}x_3 (E^{12} + E^{21}) + x_4 E^{22}, \quad (8.204)$$

where E^{ij} is defined as the $n \times n$ matrix with 1 in the $(i, j)^{\text{th}}$ coordinate and 0 elsewhere. Thus the positive semidefinite constraint can be replaced by the linear matrix inequality

$$X \succeq 0 \iff -\left(x_1 E^{11} + \frac{1}{2}x_2 (E^{12} + E^{21}) + \frac{1}{2}x_3 (E^{12} + E^{21}) + x_4 E^{22}\right) \preceq 0. \quad (8.205)$$

The general case goes the same way. We can say

$$X \succeq 0 \iff -\left(\sum_{i=1}^n x_{i,i} E^{ii} + \frac{1}{2} \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n x_{i,j} (E^{ij} + E^{ji})\right) \preceq 0 \quad (8.206)$$

where again $x_{i,j}$ refers to the element of \vec{x} corresponding to the entry X_{ij} .

This gives a linear matrix inequality for the last constraint, and so all constraints can be represented by some linear matrix inequalities. Thus, by the discussion on reducing several linear matrix inequalities to a single one, all constraints

can be represented as a single linear matrix inequality of the form $F_0 + \sum_{i=1}^m x_i F_i$. Thus the original problem can be represented as

$$\min_{\vec{x} \in \mathbb{R}^m} \vec{c}^\top \vec{x} \quad (8.207)$$

$$\text{s.t. } F_0 + \sum_{i=1}^m x_i F_i \preceq 0, \quad (8.208)$$

where $m = n^2$, $\vec{c} = \text{vec}(C)$, and the linear matrix inequality constraint is constructed in the aforementioned way. \square

Theorem 204 (Dual of an SDP)

The dual of an SDP is an SDP.

Proof. Let $\vec{c} \in \mathbb{R}^n$, and let $F_0, F_1, \dots, F_n \in \mathbb{S}^d$. Consider the following inequality-form SDP:

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.209)$$

$$\text{s.t. } F_0 + \sum_{i=1}^n x_i F_i \preceq 0. \quad (8.210)$$

We compute the conic dual of this problem. We know that the dual cone of \mathbb{S}_+^d in \mathbb{S}^d (equipped with the Frobenius inner product $\langle A, B \rangle_F = \text{tr}(AB)$ and corresponding Frobenius norm) is simply \mathbb{S}_+^d itself. Thus we can define the Lagrangian $L: \mathbb{R}^n \times \mathbb{S}_+^d$ as

$$L(\vec{x}, \Lambda) = \vec{c}^\top \vec{x} + \left\langle \Lambda, F_0 + \sum_{i=1}^n x_i F_i \right\rangle_F \quad (8.211)$$

$$= \vec{c}^\top \vec{x} + \text{tr} \left(\Lambda \left(F_0 + \sum_{i=1}^n x_i F_i \right) \right) \quad (8.212)$$

$$= \vec{c}^\top \vec{x} + \text{tr}(\Lambda F_0) + \sum_{i=1}^n x_i \text{tr}(\Lambda F_i) \quad (8.213)$$

$$= \sum_{i=1}^n (c_i + \text{tr}(\Lambda F_i)) x_i + \text{tr}(\Lambda F_0). \quad (8.214)$$

Now, define the dual function $g: \mathbb{S}_+^d \rightarrow \mathbb{R}$ by minimizing over the primal variable \vec{x} :

$$g(\Lambda) = \min_{\vec{x} \in \mathbb{R}^d} L(\vec{x}, \Lambda) \quad (8.215)$$

$$= \min_{\vec{x} \in \mathbb{R}^d} \left(\sum_{i=1}^n (c_i + \text{tr}(\Lambda F_i)) x_i + \text{tr}(\Lambda F_0) \right) \quad (8.216)$$

$$= \text{tr}(\Lambda F_0) + \sum_{i=1}^n \min_{x_i \in \mathbb{R}} (c_i + \text{tr}(\Lambda F_i)) x_i \quad (8.217)$$

$$= \begin{cases} \text{tr}(\Lambda F_0), & \text{if } \text{tr}(\Lambda F_i) = -c_i, \quad \forall i \in \{1, \dots, n\} \\ -\infty, & \text{otherwise.} \end{cases} \quad (8.218)$$

The last equality is because in each individual term $(c_i + \text{tr}(\Lambda F_i)) x_i$, when minimizing over x_i , if $c_i + \text{tr}(\Lambda F_i) \neq 0$ then we can always drive it to $-\infty$ by picking x_i to be large and of the opposite sign.

We can thus write the dual problem as

$$d^* = \max_{\Lambda \in \mathbb{S}^d} \operatorname{tr}(F_0 \Lambda) \quad (8.219)$$

$$\text{s.t. } \operatorname{tr}(F_i \Lambda) = c_i, \quad \forall i \in \{1, \dots, n\}, \quad (8.220)$$

$$\Lambda \succeq 0, \quad (8.221)$$

which is an SDP in standard form. \square

SDPs generalize all previously introduced classes of convex optimization problems: LPs, (convex) QPs, (convex) QCQPs, and SOCPs.

Theorem 205

SOCPs can be reformulated as SDPs.

Proof. We use the following useful characterization of second-order cone constraints as semidefinite constraints.

Claim. For $(\vec{x}, t) \in \mathbb{R}^{m+1}$, we have

$$\|\vec{x}\|_2 \leq t \iff \begin{bmatrix} tI & \vec{x} \\ \vec{x}^\top & t \end{bmatrix} \succeq 0. \quad (8.222)$$

Proof of claim. We have

$$\begin{bmatrix} tI & \vec{x} \\ \vec{x}^\top & t \end{bmatrix} \succeq 0 \iff \begin{bmatrix} \vec{a} \\ b \end{bmatrix}^\top \begin{bmatrix} tI & \vec{x} \\ \vec{x}^\top & t \end{bmatrix} \begin{bmatrix} \vec{a} \\ b \end{bmatrix} \geq 0, \quad \forall (\vec{a}, b) \in \mathbb{R}^{m+1} \quad (8.223)$$

$$\iff \begin{bmatrix} \vec{a} \\ b \end{bmatrix}^\top \begin{bmatrix} t\vec{a} + b\vec{x} \\ \vec{x}^\top \vec{a} + tb \end{bmatrix} \geq 0, \quad \forall (\vec{a}, b) \in \mathbb{R}^{m+1} \quad (8.224)$$

$$\iff t(\|\vec{a}\|_2^2 + b^2) + 2b\vec{a}^\top \vec{x} \geq 0, \quad \forall (\vec{a}, b) \in \mathbb{R}^{m+1}. \quad (8.225)$$

By Cauchy-Schwarz we have

$$t(\|\vec{a}\|_2^2 + b^2) + 2b\vec{a}^\top \vec{x} \geq t(\|\vec{a}\|_2^2 + b^2) - 2|b|\|\vec{a}\|_2\|\vec{x}\|_2 \quad (8.226)$$

with equality when $\vec{a} = -K\vec{x}$ for some positive scalar $K > 0$. Thus

$$\begin{bmatrix} tI & \vec{x} \\ \vec{x}^\top & t \end{bmatrix} \succeq 0 \iff t(\|\vec{a}\|_2^2 + b^2) - 2|b|\|\vec{a}\|_2\|\vec{x}\|_2 \geq 0, \quad \forall (\vec{a}, b) \in \mathbb{R}^{m+1}. \quad (8.227)$$

Now by the AM-GM inequality (i.e., expanding the square on $(\|\vec{a}\|_2 - |b|)^2 \geq 0$), we have $\|\vec{a}\|_2^2 + b^2 \geq 2|b|\|\vec{a}\|_2$, with equality when $\|\vec{a}\|_2 = |b|$. This gives

$$t(\|\vec{a}\|_2^2 + b^2) - 2|b|\|\vec{a}\|_2\|\vec{x}\|_2 \geq 2|b|\|\vec{a}\|_2(t - \|\vec{x}\|_2) \quad (8.228)$$

with equality when $\|\vec{a}\|_2 = |b|$. Thus we have

$$\begin{bmatrix} tI & \vec{x} \\ \vec{x}^\top & t \end{bmatrix} \succeq 0 \iff 2|b|\|\vec{a}\|_2(t - \|\vec{x}\|_2) \geq 0 \quad \forall (\vec{a}, b) \in \mathbb{R}^{m+1} \quad (8.229)$$

$$\iff t - \|\vec{x}\|_2 \geq 0 \quad (8.230)$$

as desired, so the claim is proved.

Now fix $\vec{c} \in \mathbb{R}^n$, fix $\vec{b}_1, \dots, \vec{b}_m \in \mathbb{R}^n$, fix A_1, \dots, A_m so that $A_i \in \mathbb{R}^{d_i \times n}$, fix $\vec{y}_1, \dots, \vec{y}_m$ so that $\vec{y}_i \in \mathbb{R}^{d_i}$, and fix $z_1, \dots, z_m \in \mathbb{R}$. Consider the following generic SOCP:

$$\min_{\vec{x} \in \mathbb{R}^n} \vec{c}^\top \vec{x} \quad (8.231)$$

$$\text{s.t.} \quad \|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i, \quad \forall i \in \{1, \dots, m\}. \quad (8.232)$$

Let K^d be the second-order cone in \mathbb{R}^d . Notice that each cone constraint can be written in the form

$$\|A_i \vec{x} - \vec{y}_i\|_2 \leq \vec{b}_i^\top \vec{x} + z_i \quad (8.233)$$

$$\iff (A_i \vec{x} - \vec{y}_i, \vec{b}_i^\top \vec{x} + z_i) \in K^{d_i+1} \quad (8.234)$$

$$\iff \begin{bmatrix} (\vec{b}_i^\top \vec{x} + z_i)I & A_i \vec{x} - \vec{y}_i \\ (A_i \vec{x} - \vec{y}_i)^\top & \vec{b}_i^\top \vec{x} + z_i \end{bmatrix} \succeq 0 \quad (8.235)$$

$$\iff \begin{bmatrix} z_i I & -\vec{y}_i \\ -\vec{y}_i^\top & z_i \end{bmatrix} + \sum_{j=1}^n x_j \begin{bmatrix} (\vec{b}_i)_j I & (A_i)_j \\ (A_i)_j^\top & (\vec{b}_i)_j \end{bmatrix} \succeq 0 \quad (8.236)$$

where $(\vec{b}_i)_j$ is the j^{th} entry of \vec{b}_i , and $(A_i)_j$ is the j^{th} column of A_i . Anyways, this is a linear matrix inequality (after some reshuffling of terms).

The conversion from SOCP to SDP consists of converting all second-order cone constraints to small linear matrix inequalities, then combining them to form one larger linear matrix inequality, which defines the constraint set of the inequality-form SDP. The objective function is already linear, so the resulting SDP is in the “standard” inequality form. Thus, we have reduced the original SOCP to an SDP. \square

Note that in practice, this reduction is often extremely costly; SDPs are hard to solve at large scale, while SOCPs are much easier.

The above content is optional/out of scope for this semester, but now we resume the required/in scope content.

8.6 General Taxonomy

We conclude this chapter with a taxonomy of problems that we have discussed until now:

$$\text{LPs} \subset \text{Convex QPs} \subset \text{Convex QCQPs} \subset \text{SOCPs} \subset \text{SDPs} \subset \text{Convex Problems} \quad (8.237)$$

All inclusions are strict, i.e., none of the classes is equivalent to any of the others.

For extra optional reading, you may also look into *geometric programs (GPs)*, which are nonconvex programs that can be turned into convex programs with a change of variables; and *mixed-integer programs (MIPs)*, which are useful in practice to incorporate integer constraints, but difficult to solve exactly. All such material is out of scope of the course.

Chapter 9

Regularization and Sparsity

Relevant sections of the textbooks:

- [2] Chapters 9, 12, 13.

9.1 Recapping Ridge Regression and Defining LASSO

The first example of regularization we saw was ridge regression. In this section, we'll review ridge regression. The most basic perspective of ridge regression focuses on the additional term we add to the objective function. In ridge regression, we solve the following problem:

$$\min_{\vec{x} \in \mathbb{R}^n} \left\{ \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_2^2 \right\}. \quad (9.1)$$

This is different from the OLS problem due to the additional $\lambda \|\vec{x}\|_2^2$ term, which can be thought of as a *regularizer* (i.e., a penalty) for having large \vec{x} values. The λ parameter controls the strength of the penalty and is usually called a *regularization parameter*. In this sense, ridge regression is *regularized least squares*. More generally, we may define regularization as follows.

Definition 206 (Regularization)

Consider the optimization problem

$$p^* = \min_{\vec{x} \in \Omega} f_0(\vec{x}). \quad (9.2)$$

For a given function $R: \Omega \rightarrow \mathbb{R}_+$ (the *regularizer*) and a *regularization parameter* $\lambda > 0$, the *regularized version* of the above problem is the problem

$$p_\lambda^* = \min_{\vec{x} \in \Omega} \{f_0(\vec{x}) + \lambda R(\vec{x})\}. \quad (9.3)$$

Here λ controls the strength of the regularization.

In general, the original problem and the regularized problem *do not have the same solutions*, nor do versions of the regularized problem with different λ parameter. One need only consider ridge regression to keep this in mind; for a fixed A and \vec{y} , increasing λ will decrease the norm of the solution to the ridge regression problem, and sending it to 0 (i.e., recovering unregularized least squares) will increase the norm of the solution.

One example of regularization is the ℓ^2 -norm penalty $R(\vec{x}) = \|\vec{x}\|_2^2$, which (when combined with $f_0(\vec{x}) = \|A\vec{x} - \vec{y}\|_2^2$) yields ridge regression. Another example is the elastic-net regression, which we covered briefly as an

example when discussing convexity. But the main objective of this chapter is to look at the so-called LASSO regression problem, which uses an ℓ^1 -norm regularizer. Recall that for a vector $\vec{x} \in \mathbb{R}^n$, its ℓ^1 -norm is defined as $\|\vec{x}\|_1 = \sum_{i=1}^n |x_i|$.

Definition 207 (LASSO Regression)

Let $A \in \mathbb{R}^{m \times n}$, $\vec{y} \in \mathbb{R}^m$, and $\lambda > 0$. The *LASSO regression* problem is:

$$\min_{\vec{x} \in \mathbb{R}^n} \left\{ \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_1 \right\}. \quad (9.4)$$

Here are some key properties of the LASSO regression problem.

Proposition 208

Consider the LASSO regression problem

$$\min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad \text{where} \quad f_0(\vec{x}) \doteq \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_1. \quad (9.5)$$

- (a) The function $f_0: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex.
- (b) If A has full column rank then f_0 is μ -strongly convex with $\mu = 2\sigma_n\{A\}^2$.
- (c) A solution $\vec{x}^* \in \operatorname{argmin}_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x})$ always exists.
- (d) If A has full column rank then the above solution is unique.

This picture is *very* different from ridge regression, where we are guaranteed that a solution always exists, is unique, and solvable in closed form. The question then becomes: why do we even care about the LASSO problem at all? The basic answer is that *it induces sparsity in the solution*, i.e., solutions to LASSO usually tend to have few nonzero entries. This sparsity is useful for applications in high-dimensional statistics and machine learning, as it reveals a certain structure — in words, it points out which “features” are the most relevant to the regression. In the following sections, we will observe how this sparsity emerges, both geometrically and algebraically.

9.2 Understanding the Difference Between the ℓ^2 -Norm and the ℓ^1 -Norm

In this section, we attempt to build more intuition about the difference between the ℓ^2 -norm and the ℓ^1 -norm. We do this by solving some problems which use the ℓ^2 norm, then replace it with the ℓ^1 norm and solve this new problem. Besides giving us intuition, it will help us learn how to analyze the LASSO problem.

Here is a diagram of the norm balls of the ℓ^1 (blue) and ℓ^2 (red) norms in $n = 1$ dimensions:




Figure 9.1: The ℓ^1 and ℓ^2 norm balls in $n = 2$ dimensions. Recall that the ℓ^p -norm ball is defined as the set of vectors \vec{v} such that $\|\vec{v}\|_p \leq 1$.

The border of the norm balls are the points where each norm is equal to 1. Notice the difference in the geometry of these norm balls. The ℓ^2 norm ball is circular, while the ℓ^1 norm ball has distinctive corners.

In fact, these corners hint at a key difference between these norms: the ℓ^2 norm is differentiable everywhere, but $\|\vec{x}\|_1$ is *not differentiable* when *any* $x_i = 0$. These corners will help us understand how the ℓ^1 norm regularizer induces sparsity in the solution, and also inform our analysis of problems involving the ℓ^1 -norm, including LASSO.

Example 209 (Least ℓ^1 -Norm). Recall that we solved the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \|\vec{x}\|_2^2 \tag{9.6}$$

$$\text{s.t. } A\vec{x} = \vec{y}. \tag{9.7}$$

Using the KKT conditions, namely stationarity, we found an explicit solution to this problem: $\vec{x}^* = A^\top(AA^\top)^{-1}\vec{y}$. Now let us replace the ℓ^2 norm with an ℓ^1 norm; we obtain the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \|\vec{x}\|_1 \tag{9.8}$$

$$\text{s.t. } A\vec{x} = \vec{y}. \tag{9.9}$$

We cannot apply stationarity to this problem because the objective is non-differentiable. Thus, this problem seems intractable to solve by hand, at least for the moment. Instead, let us formulate it as a linear program. As before, we represent each x_i as the difference of non-negative numbers which sum to $|x_i|$. More formally, we introduce slack variables $\vec{x}^+, \vec{x}^- \in \mathbb{R}^n$ such that for each $i \in \{1, \dots, n\}$ we have $x_i^+ \geq 0, x_i^- \geq 0, x_i^+ - x_i^- = x_i$, and $x_i^+ + x_i^- = |x_i|$. Thus we can rewrite the problem using the following linear program:

$$\min_{\vec{x}^+, \vec{x}^- \in \mathbb{R}^n} \sum_{i=1}^n (x_i^+ + x_i^-) \tag{9.10}$$

$$\text{s.t. } A(\vec{x}^+ - \vec{x}^-) = \vec{y} \tag{9.11}$$

$$\vec{x}^+ \geq \vec{0} \tag{9.12}$$

$$\vec{x}^- \geq \vec{0}. \tag{9.13}$$

This is a linear program which is efficiently solvable.

As a corollary, we can consider the ℓ^1 -norm regression problem:

$$\min_{\vec{x} \in \mathbb{R}^n} \|A\vec{x} - \vec{y}\|_1. \tag{9.14}$$

We can introduce the slack variable $\vec{e} = A\vec{x} - \vec{y}$ and obtain the problem:

$$\min_{\substack{\vec{x} \in \mathbb{R}^n \\ \vec{e} \in \mathbb{R}^m}} \|\vec{e}\|_1 \tag{9.15}$$

$$\text{s.t. } A\vec{x} - \vec{y} = \vec{e} \tag{9.16}$$

which is an equality-constrained ℓ^1 minimization problem, and thus a linear program as demonstrated above.

Example 210 (Mean Versus Median). Let k be a positive integer. Suppose we have points $\vec{x}_1, \dots, \vec{x}_k \in \mathbb{R}^n$. Consider the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^k \|\vec{x} - \vec{x}_i\|_2^2. \tag{9.17}$$

This is an unconstrained strongly convex differentiable problem, so it has a unique solution \vec{x}_1^* which we may find by setting the derivative of the objective to $\vec{0}$. We obtain

$$\vec{0} = 2 \sum_{i=1}^k (\vec{x}_1^* - \vec{x}_i) \tag{9.18}$$

$$\implies \vec{0} = \sum_{i=1}^k (\vec{x}_1^* - \vec{x}_i) = k \cdot \vec{x}_1^* - \sum_{i=1}^k \vec{x}_i \tag{9.19}$$

$$\implies \vec{x}_1^* = \frac{1}{k} \sum_{i=1}^k \vec{x}_i. \tag{9.20}$$

This computation implies that the sample mean is the point which minimizes the total squared distance to all points in the dataset.

Now suppose that we instead consider the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \sum_{i=1}^k \|\vec{x} - \vec{x}_i\|_2. \tag{9.21}$$

The solution to this problem is the *sample median* of the points. To see this, suppose that $n = 1$, i.e., all our data x_i are scalar-valued. Then we obtain the problem

$$\min_{x \in \mathbb{R}} \sum_{i=1}^k |x - x_i|. \tag{9.22}$$

This is an unconstrained, convex, non-differentiable problem. Let us examine all critical points – that is, points where the derivative is 0 or undefined. The derivative of the objective is

$$\frac{d}{dx} \sum_{i=1}^k |x - x_i| = \sum_{i=1}^k \frac{d}{dx} |x - x_i| \tag{9.23}$$

$$= \sum_{i=1}^k \begin{cases} 1, & \text{if } x > x_i \\ -1, & \text{if } x < x_i \\ \text{undefined,} & \text{if } x = x_i \end{cases} \tag{9.24}$$

$$= \begin{cases} \sum_{i: x > x_i} 1 + \sum_{i: x < x_i} -1, & \text{if } x \notin \{x_1, \dots, x_k\} \\ \text{undefined,} & \text{if } x \in \{x_1, \dots, x_k\} \end{cases} \tag{9.25}$$

$$= \begin{cases} |\{i \in \{1, \dots, k\} : x > x_i\}| - |\{i \in \{1, \dots, k\} : x < x_i\}|, & \text{if } x \notin \{x_1, \dots, x_k\} \\ \text{undefined}, & \text{if } x \in \{x_1, \dots, x_k\}. \end{cases} \quad (9.26)$$

Thus if x is such that $|\{i \in \{1, \dots, k\} : x > x_i\}| = |\{i \in \{1, \dots, k\} : x < x_i\}|$, then the derivative is 0, so this x is a candidate solution. To put this convoluted-looking condition in words, notice that the first term in the equality is just the number of x_i which are larger than x , and the second term is the number of x_i which are smaller than x . Thus the condition says that there are the same number of points in the set which are larger than x as there are points which are smaller than x . This x would fulfill the traditional definition of “median” as the middle of the sorted list of points.

To formally solve this problem, one must also check all the values $x = x_i$ and compare the objective values. But eventually after doing all this, one recovers that the optimal solutions are all possible medians of the dataset.

Because the median is defined using the $|\cdot|$ instead of $(\cdot)^2$ function, it inherits several different properties. The most striking is its robustness; the median is much more robust than the mean. The mean is very sensitive to outliers, while the median is less sensitive (i.e. if we blow up an outlier point, the mean will change a lot, while the median will be unaffected).

9.3 Analysis of LASSO Regression

In this section we will solve the one-dimensional LASSO problem. The ideas generalize to the vector case directly, through a reduction of the vector LASSO problems to several one-dimensional LASSO problems. The details of this reduction are left as a homework exercise.

First consider the scalar ridge regression problem, with $\vec{a} \neq \vec{0}$:

$$\min_{x \in \mathbb{R}} f_{\text{RR}}(x) \quad \text{where} \quad f_{\text{RR}}(x) \doteq \frac{1}{2} \|\vec{a}x - \vec{y}\|_2^2 + \frac{1}{2} \lambda x^2. \quad (9.27)$$

By taking the derivative, we get

$$\frac{df_{\text{RR}}}{dx}(x) = \vec{a}^\top (\vec{a}x - \vec{y}) + \lambda x \quad (9.28)$$

$$= (\vec{a}^\top \vec{a} + \lambda)x - \vec{a}^\top \vec{y} \quad (9.29)$$

$$= (\|\vec{a}\|_2^2 + \lambda)x - \vec{a}^\top \vec{y} \quad (9.30)$$

$$\implies x_{\text{RR}}^* = \frac{\vec{a}^\top \vec{y}}{\|\vec{a}\|_2^2 + \lambda}. \quad (9.31)$$

By setting $\lambda = 0$ we obtain the least squares solution:

$$x_{\text{LS}}^* = \frac{\vec{a}^\top \vec{y}}{\|\vec{a}\|_2^2} \quad (9.32)$$

Now consider the scalar LASSO problem

$$\min_{x \in \mathbb{R}} f_{\text{LASSO}}(x) \quad \text{where} \quad f_{\text{LASSO}}(x) \doteq \frac{1}{2} \|\vec{a}x - \vec{y}\|_2^2 + \lambda |x|. \quad (9.33)$$

This has a derivative everywhere except $x = 0$. We obtain

$$\frac{df_{\text{LASSO}}}{dx}(x) = \vec{a}^\top (\vec{a}x - \vec{y}) + \lambda \begin{cases} 1, & \text{if } x > 0 \\ -1, & \text{if } x < 0 \\ \text{undefined}, & \text{if } x = 0. \end{cases} \quad (9.34)$$

Let x^* be a critical point of this problem. We solve what x^* should be using casework.

Case 1. If $x^* > 0$, then the derivative is well-defined, so it must be equal to 0. Thus we have

$$0 = \frac{df_{\text{LASSO}}}{dx}(x^*) \quad (9.35)$$

$$= \vec{a}^\top (\vec{a}x^* - \vec{y}) + \lambda \quad (9.36)$$

$$\implies x^* = \frac{\vec{a}^\top \vec{y} - \lambda}{\|\vec{a}\|_2^2}. \quad (9.37)$$

Thus if $x^* > 0$ then $\vec{x}^* = \frac{\vec{a}^\top \vec{y} - \lambda}{\|\vec{a}\|_2^2}$. Thus $\vec{x}^* > 0$ if and only if $\vec{a}^\top \vec{y} > \lambda$.

Case 2. If $x^* < 0$, then the derivative is well-defined, so it must be equal to 0. Thus we have

$$0 = \frac{df_{\text{LASSO}}}{dx}(x^*) \quad (9.38)$$

$$= \vec{a}^\top (\vec{a}x^* - \vec{y}) - \lambda \quad (9.39)$$

$$\implies x^* = \frac{\vec{a}^\top \vec{y} + \lambda}{\|\vec{a}\|_2^2}. \quad (9.40)$$

Thus if $x^* < 0$ then $\vec{x}^* = \frac{\vec{a}^\top \vec{y} + \lambda}{\|\vec{a}\|_2^2}$. Thus $\vec{x}^* < 0$ if and only if $\vec{a}^\top \vec{y} < -\lambda$.

Case 3. If $x^* = 0$ then it is neither > 0 nor < 0 . Thus we must have $-\lambda \leq \vec{a}^\top \vec{y} \leq \lambda$.

Thus we have three cases:

- $x^* > 0 \iff \vec{a}^\top \vec{y} > \lambda$, in which case $x^* = (\vec{a}^\top \vec{y} - \lambda) / \|\vec{a}\|_2^2$;
- $x^* < 0 \iff \vec{a}^\top \vec{y} < -\lambda$, in which case $x^* = (\vec{a}^\top \vec{y} + \lambda) / \|\vec{a}\|_2^2$; and
- $x^* = 0 \iff -\lambda \leq \vec{a}^\top \vec{y} \leq \lambda$,

or in other words,

$$x^* = \begin{cases} (\vec{a}^\top \vec{y} - \lambda) / \|\vec{a}\|_2^2, & \text{if } \vec{a}^\top \vec{y} > \lambda \\ (\vec{a}^\top \vec{y} + \lambda) / \|\vec{a}\|_2^2, & \text{if } \vec{a}^\top \vec{y} < -\lambda \\ 0, & \text{if } -\lambda \leq \vec{a}^\top \vec{y} \leq \lambda. \end{cases} \quad (9.41)$$

As a function of $\vec{a}^\top \vec{y}$, the solution $x^* = x_{\text{LASSO}}^*$ looks like:




Figure 9.2: The plot of the function which maps $\vec{a}^\top \vec{y} \mapsto x_{\text{LASSO}}^*$, where the latter term is the solution to our scalar LASSO problem. When $\vec{a}^\top \vec{y} \in [-\lambda, \lambda]$, we have $x^* = 0$. The function $\vec{a}^\top \vec{y} \mapsto \vec{x}^*$ is continuous, yet not differentiable at $\vec{a}^\top \vec{y} = \pm\lambda$.

If we plot the least squares solution in red on the same graph, it has the same nonzero slope as the LASSO solution, and looks like this:




Figure 9.3: In red, we add the plot of the function which maps $\bar{a}^\top \bar{y} \mapsto x_{\text{LS}}^*$, where the latter term is the solution to our scalar least squares problem. Note that the LASSO solution (in blue) is always closer to zero than the least squares solution, and is set directly to zero when $\bar{a}^\top \bar{y} \in [-\lambda, \lambda]$.

This illustrates a concept called *soft thresholding*: in the regime where the least squares solution x_{LS}^* is already close to zero, x_{LASSO}^* becomes exactly zero. Meanwhile, ridge regression does not do this: $x_{\text{RR}}^* = 0$ if and only if $\bar{a}^\top \bar{y} = 0$, which is exactly when the unregularized least squares solution itself is zero. This fundamental difference is why the solutions to LASSO regression tend to be sparse, i.e., have many entries set to 0.

9.4 Geometry of LASSO Regression

In this section we introduce a geometric description of LASSO. The geometry relies on a crucial theorem which unveils a deep connection between regularization and constrained optimization for convex problems. We state the result here; the proof uses duality theory and is left to homework.

Theorem 211

Let $f_0: \mathbb{R}^n \rightarrow \mathbb{R}$ be strictly convex and such that $\lim_{\alpha \rightarrow \infty} f_0(\alpha \vec{x}) = \infty$ for all $\vec{x} \neq \vec{0}$,^a and $R: \mathbb{R}^n \rightarrow \mathbb{R}_+$ be convex and take non-negative values. Further suppose that there exists $\vec{x}_0 \in \mathbb{R}^n$ such that $R(\vec{x}_0) = 0$.

For $\lambda \geq 0$ and $t \geq 0$, let $\mathcal{R}(\lambda)$ and $\mathcal{C}(t)$ be sets of solutions to the “regularized” and “constraint” programs:

$$\mathcal{R}(\lambda) \doteq \underset{\vec{x} \in \mathbb{R}^n}{\operatorname{argmin}} \{f_0(\vec{x}) + \lambda R(\vec{x})\} \quad (9.42)$$

$$\mathcal{C}(t) \doteq \underset{\substack{\vec{x} \in \mathbb{R}^n \\ R(\vec{x}) \leq t}}{\operatorname{argmin}} f_0(\vec{x}). \quad (9.43)$$

Then:

- (a) for every $\lambda \geq 0$ there exists $t \geq 0$ such that $\mathcal{R}(\lambda) = \mathcal{C}(t)$; and
- (b) for every $t > 0$ there exists $\lambda \geq 0$ such that $\mathcal{R}(\lambda) = \mathcal{C}(t)$.

^aThis assumption is called “coercivity”.

This shows that in some sense, regularized convex problems are equivalent to constrained convex problems; and in this equivalence, the regularizer for the regularized problem shapes the constraint set of the constrained problem. In particular, regularized least squares ($f_0(\vec{x}) = \|A\vec{x} - \vec{y}\|_2^2$) with full column rank is equivalent to constrained least squares (with the same f_0).

Now, we sketch the feasible sets and level sets of the objective function for the constrained problems corresponding to both ridge regression and LASSO regression.




Figure 9.4: Geometric differences between LASSO and ridge regression. On the left side, the blue diamond depicts the feasible region for an ℓ^1 -norm constraint such as $\|\vec{x}\|_1 \leq t$, while the circle on the right side is the feasible region for an ℓ^2 -norm constraint such as $\|\vec{x}\|_2^2 \leq t$. On both graphs, the red line is a level set of our objective function; specifically, the minimal level set that still intersects the feasible region. The intersection of this level set with the feasible region is the solution to our constrained problem and thus to an equivalent regularized problem.

Note how with the ℓ^1 -norm constraint, the intersection of the feasible region with the minimal level set is more likely to be at a corner of the feasible region, which is a point where some coordinates are set exactly to zero. Meanwhile, with the ℓ^2 -norm constraint, the intersection can be at an arbitrary point on the circle (or sphere in higher dimensions), and likely isn't at a corner. This is why LASSO induces sparsity in \vec{x} , due to the distinctive corners we saw earlier in its norm ball. Meanwhile, although ridge regression compresses \vec{x}_{RR}^* to be smaller, it doesn't necessarily induce sparsity in \vec{x}_{RR}^* .

Chapter 10

Advanced Descent Methods

Relevant sections of the textbooks:

- [1] Chapter 11.
- [2] Section 12.5.

10.1 Coordinate Descent

Recall from Chapter 6 the general idea of descent-based methods for solving optimization problems. These methods are based on the process of starting with some initial guess $\vec{x}^{(0)} \in \mathbb{R}^n$, then generating a sequence of refined guesses $\vec{x}^{(1)}, \vec{x}^{(2)}, \vec{x}^{(3)}, \dots$ using the general update rule

$$\vec{x}^{(t+1)} = \vec{x}^{(t)} + \eta \vec{v}^{(t)} \quad (10.1)$$

for some search direction $\vec{v}^{(t)}$ and step size η . In Chapter 6 we covered the gradient descent method, which uses the gradient of the function as the search direction. In this chapter we will revisit descent-based optimization methods and introduce alternative update rules.

In this section we will introduce *coordinate descent*, a class of descent-based algorithms that finds a minimizer of multivariate functions by iteratively minimizing it along one direction at a time. Consider the unconstrained convex optimization problem

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}), \quad (10.2)$$

with the optimization variable

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}. \quad (10.3)$$

Before we introduce the algorithm, we introduce some notation. For indices i and j , we introduce the notation $\vec{x}_{i:j} = (x_i, x_{i+1}, \dots, x_{j-1}, x_j) \in \mathbb{R}^{j-i+1}$ to be the entries of \vec{x} between indices i and j (inclusive on both ends).

Given an initial guess $\vec{x}^{(0)}$, for $t \geq 0$ the coordinate descent algorithm updates the iterate $\vec{x}^{(t)}$ by *sequentially* minimizing the function $f(\vec{x})$ with respect to each coordinate, namely using the update rule

$$x_i^{(t+1)} \in \operatorname{argmin}_{x_i \in \mathbb{R}} f(\vec{x}_{1:i-1}^{(t+1)}, x_i, \vec{x}_{i+1:n}^{(t)}). \quad (10.4)$$

Namely, we perform the update

$$x_1^{(t+1)} \in \operatorname{argmin}_{x_1 \in \mathbb{R}} f(x_1, \vec{x}_{2:n}^{(t)}) \quad (10.5)$$

$$x_2^{(t+1)} \in \operatorname{argmin}_{x_2 \in \mathbb{R}} f(x_1^{(t+1)}, x_2, \vec{x}_{3:n}^{(t)}) \quad (10.6)$$

$$\vdots \quad \quad \quad \vdots \quad (10.7)$$

$$x_n^{(t+1)} \in \operatorname{argmin}_{x_n \in \mathbb{R}} f(\vec{x}_{1:n-1}^{(t+1)}, x_n). \quad (10.8)$$

This is a sequential process since after finding the minimizer along the i^{th} coordinate (i.e. $x_i^{(t+1)}$) we use its values for minimizing subsequent coordinates. Also we note that the order of the coordinates is arbitrary. We formalize this update in the following algorithm.

Algorithm 6 CoordinateDescent

```

1: function COORDINATEDESCENT( $f, \vec{x}^{(0)}, T$ )
2:   for  $t = 0, 1, \dots, T_1$  do
3:     for  $i = 1, \dots, N$  do
4:        $x_i^{(t+1)} \leftarrow \operatorname{argmin}_{x_i \in \mathbb{R}} f(x_{1:i-1}^{(t+1)}, x_i, x_{i+1:n}^{(t)})$ .
5:     end for
6:   end for
7:   return  $\vec{x}^T$ 
8: end function

```

The algorithm breaks down the difficult multivariate optimization problem into a sequence of simpler univariate optimization problems.

We first want to discuss the issue of well-posedness of the algorithm. We know that any of the argmins used may not exist, in which case the algorithm is not well-defined, and so we cannot even think about its behavior or convergence. Nevertheless, in a large class of problems which have many different characterizations, the argmins are well-defined. We say in this case that the coordinate descent algorithm is well-posed.

We now want to address the question of convergence. It is not obvious that minimizing the function $f(\vec{x})$ can be achieved by minimizing along each direction separately. In fact, the algorithm is not guaranteed to converge to an optimal solution for general convex functions. However, under some additional assumptions on the function, we can guarantee convergence. To build an intuition for what additional assumptions are needed we consider the following question. Let $f(\vec{x})$ be a convex differentiable function. Suppose that $x_i^* \in \operatorname{argmin}_{x_i \in \mathbb{R}} f(x_{1:i-1}^*, x_i, x_{i+1:n}^*)$ for all i . Can we conclude that \vec{x}^* is a global minimizer of $f(\vec{x})$? The answer to this question is yes. We can prove this by recalling the first order optimality condition for unconstrained convex functions and the definition of partial derivatives. If \vec{x}^* is a minimizer of $f(\vec{x})$ along the direction \vec{e}_i then we have

$$\frac{\partial f}{\partial x_i}(\vec{x}^*) = 0. \quad (10.9)$$

If this is true for all i then $\nabla f(\vec{x}^*) = \vec{0}$, implying that \vec{x}^* is a global minimizer for f . This discussion forms a proof of the following theorem, which is Theorem 12.4 in [2].

Theorem 212 (Convergence of Coordinate Descent for Differentiable Convex Functions)

Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable convex function which is *separately strictly convex* in each argument. That is, suppose that for each i , and each fixed $\vec{x}_{1:i-1}$ and $\vec{x}_{i+1:n}$, the function $x_i \mapsto f(\vec{x}_{1:i-1}, x_i, \vec{x}_{i+1:n})$ is strictly convex. If the coordinate descent algorithm is well-posed, and the unconstrained minimization problem

$$\min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}) \quad (10.10)$$

has a solution, then the sequence of iterates $\vec{x}^{(0)}, \vec{x}^{(1)}, \dots$ generated by the coordinate descent algorithm converges to an optimal solution to (10.10).

The coordinate descent algorithm may not converge to an optimal solution for general non-differentiable functions, even if they are convex. However, we can still prove that coordinate descent converges for a special class of functions of the form

$$f(\vec{x}) = g(\vec{x}) + \sum_{i=1}^n h_i(x_i) \quad (10.11)$$

where $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and differentiable, and each $h_i: \mathbb{R} \rightarrow \mathbb{R}$ is convex (but not necessarily differentiable). This form includes various ℓ^1 regularization problems (such as LASSO regression) which have a separable non-differentiable component. The provable convergence of coordinate descent algorithm makes it an attractive choice for this class of problems.

Example 213. In this example we will consider the LASSO regression problem and examine how coordinate descent algorithm can be applied to solve it. Note that the LASSO objective follows the form described in (10.11). For $A \in \mathbb{R}^{m \times n}$ which has columns $\vec{a}_1, \dots, \vec{a}_n \in \mathbb{R}^m$, and $\vec{y} \in \mathbb{R}^m$, we consider the LASSO objective

$$f(\vec{x}) = \frac{1}{2} \|A\vec{x} - \vec{y}\|_2^2 + \lambda \|\vec{x}\|_1. \quad (10.12)$$

We aim to use coordinate descent to minimize this function. Let $\vec{x}^{(0)}$ be the initial guess. Then we perform the coordinate descent update by solving the following optimization problem:

$$x_i^{(t+1)} = \operatorname{argmin}_{x_i \in \mathbb{R}} f(\vec{x}_{1:i-1}^{(t+1)}, x_i, \vec{x}_{i+1:n}^{(t)}). \quad (10.13)$$

Each of these optimization problems will be solved similarly to what we did in Section 9.3. For notational clarity, let us instead solve the more generic problem

$$x_i^* \in \operatorname{argmin}_{x_i \in \mathbb{R}} f(\vec{x}_{1:i-1}, x_i, \vec{x}_{i+1:n}). \quad (10.14)$$

Then we have that $\frac{\partial f}{\partial x_i}(\vec{x}_{1:i-1}, x_i^*, \vec{x}_{i+1:n})$ is either 0 or undefined. Thus we compute each partial derivative of f . Indeed, we have

$$\frac{\partial f}{\partial x_i}(\vec{x}) = \frac{1}{2} \frac{\partial}{\partial x_i} \|A\vec{x} - \vec{y}\|_2^2 + \lambda \frac{\partial}{\partial x_i} \|\vec{x}\|_1 \quad (10.15)$$

$$= \vec{e}_i^\top \nabla \left(\frac{1}{2} \|A\vec{x} - \vec{y}\|_2^2 \right) + \lambda \frac{\partial}{\partial x_i} \sum_{j=1}^n |x_j| \quad (10.16)$$

$$= \vec{e}_i^\top A^\top (A\vec{x} - \vec{y}) + \lambda \frac{\partial}{\partial x_i} |x_i| \quad (10.17)$$

$$= \vec{a}_i^\top (A\vec{x} - \vec{y}) + \lambda \begin{cases} \operatorname{sgn}(x_i), & \text{if } x_i \neq 0 \\ \text{undefined}, & \text{if } x_i = 0. \end{cases} \quad (10.18)$$

We now introduce the additional notation $A_{i:j} \in \mathbb{R}^{m \times (j-i+1)}$ as the sub-matrix of A whose columns are the i^{th} through j^{th} columns of A (inclusive). Using this notation, we can simplify the first term as

$$\vec{a}_i^\top (A\vec{x} - \vec{y}) = \vec{a}_i^\top A\vec{x} - \vec{a}_i^\top \vec{y} \quad (10.19)$$

$$= \vec{a}_i^\top \left(\sum_{j=1}^n \vec{a}_j x_j - \vec{y} \right) \quad (10.20)$$

$$= \|\vec{a}_i\|_2^2 x_i + \vec{a}_i^\top \left(\sum_{\substack{j=1 \\ j \neq i}}^n \vec{a}_j x_j - \vec{y} \right) \quad (10.21)$$

$$= \|\vec{a}_i\|_2^2 + \vec{a}_i^\top (A_{1:i-1}\vec{x}_{1:i-1} + A_{i+1:n}\vec{x}_{i+1:n} - \vec{y}) \quad (10.22)$$

Now there are two cases depending on the value of x_i^* .

Case 1. Suppose $x_i^* \neq 0$. Then we have by the optimality condition that

$$0 = \frac{\partial f}{\partial x_i}(\vec{x}_{1:i-1}, x_i^*, \vec{x}_{i+1:n}) \quad (10.23)$$

$$= \|\vec{a}_i\|_2^2 x_i^* + \vec{a}_i^\top (A_{1:i-1}\vec{x}_{1:i-1} + A_{i+1:n}\vec{x}_{i+1:n} - \vec{y}) + \lambda \operatorname{sgn}(x_i^*) \quad (10.24)$$

$$\implies x_i^* = \frac{1}{\|\vec{a}_i\|_2^2} (\vec{a}_i^\top [\vec{y} - A_{1:i-1}\vec{x}_{1:i-1} - A_{i+1:n}\vec{x}_{i+1:n}] - \lambda \operatorname{sgn}(x_i^*)). \quad (10.25)$$

Thus in particular we have

$$x_i^* > 0 \iff \vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1} - A_{i+1:n}\vec{x}_{i+1:n}) > \lambda, \quad (10.26)$$

and the corresponding relation

$$x_i^* < 0 \iff \vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1} - A_{i+1:n}\vec{x}_{i+1:n}) < -\lambda. \quad (10.27)$$

Case 2. By contrapositive, we have

$$-\lambda \leq \vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1} - A_{i+1:n}\vec{x}_{i+1:n}) \leq \lambda \iff x_i^* = 0. \quad (10.28)$$

Therefore we have derived a closed-form coordinate descent update:

$$\vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1}^{(t+1)} - A_{i+1:n}\vec{x}_{i+1:n}^{(t)}) > \lambda \implies x_i^{(t+1)} = \frac{1}{\|\vec{a}_i\|_2^2} (\vec{a}_i^\top [\vec{y} - A_{1:i-1}\vec{x}_{1:i-1}^{(t+1)} - A_{i+1:n}\vec{x}_{i+1:n}^{(t)}] - \lambda) \quad (10.29)$$

$$\vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1}^{(t+1)} - A_{i+1:n}\vec{x}_{i+1:n}^{(t)}) < -\lambda \implies x_i^{(t+1)} = \frac{1}{\|\vec{a}_i\|_2^2} (\vec{a}_i^\top [\vec{y} - A_{1:i-1}\vec{x}_{1:i-1}^{(t+1)} - A_{i+1:n}\vec{x}_{i+1:n}^{(t)}] + \lambda) \quad (10.30)$$

$$\left| \vec{a}_i^\top (\vec{y} - A_{1:i-1}\vec{x}_{1:i-1}^{(t+1)} - A_{i+1:n}\vec{x}_{i+1:n}^{(t)}) \right| \leq \lambda \implies x_i^{(t+1)} = 0. \quad (10.31)$$

10.2 Newton's Method

Let us consider the unconstrained optimization problem

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}). \quad (10.32)$$

Recall that in the gradient descent algorithm, we assumed that the objective function $f(\vec{x})$ is differentiable. Furthermore, we assumed that at every point $\vec{x} \in \mathbb{R}^n$ we can compute $f(\vec{x})$ as well as $\nabla f(\vec{x})$. Here, we make the additional assumption that $f(\vec{x})$ is twice differentiable and that we can compute the Hessian $\nabla^2 f(\vec{x})$. We wish to use the Hessian to choose a good search direction and accelerate convergence. Optimization algorithms that utilize second derivatives (e.g. the Hessian) are called *second-order methods*.

One of the most famous second-order methods is *Newton's method*. Newton's method is based on the following idea for minimizing strictly-convex functions with positive definite Hessians: first, start with an initial guess $\vec{x}^{(0)}$. Then in each iteration $t = 1, 2, 3, \dots$, approximate the objective function with its second-order Taylor approximation around the point $\vec{x}^{(t)}$. The minimizer of this quadratic approximation is then chosen as the next iterate $\vec{x}^{(t+1)}$.

More formally, let us assume that f is strictly convex and twice-differentiable with positive definite Hessian at $\vec{x}^{(t)}$, and let us write the second-order Taylor approximation of the function $f(\vec{x})$ around the point $\vec{x}^{(t)}$. We obtain

$$\widehat{f}_2(\vec{x}; \vec{x}^{(t)}) = f(\vec{x}^{(t)}) + [\nabla f(\vec{x}^{(t)})]^\top (\vec{x} - \vec{x}^{(t)}) + \frac{1}{2} (\vec{x} - \vec{x}^{(t)})^\top [\nabla^2 f(\vec{x}^{(t)})] (\vec{x} - \vec{x}^{(t)}). \quad (10.33)$$

Since the Hessian $\nabla^2 f(\vec{x}^{(t)}) \succ 0$, we can solve the problem

$$\min_{\vec{x} \in \mathbb{R}^n} \widehat{f}_2(\vec{x}; \vec{x}^{(t)}), \quad (10.34)$$

which is a convex quadratic program, using our (by now) standard techniques. Setting the gradient (in \vec{x}) to $\vec{0}$, we obtain

$$\vec{0} = \nabla \widehat{f}_2(\vec{x}^*; \vec{x}^{(t)}) \quad (10.35)$$

$$= \nabla f(\vec{x}^{(t)}) + [\nabla^2 f(\vec{x}^{(t)})] (\vec{x}^* - \vec{x}^{(t)}) \quad (10.36)$$

$$\implies \vec{x}^* = \vec{x}^{(t)} - [\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]. \quad (10.37)$$

This gives the Newton's method update rule:

$$\vec{x}^{(t+1)} = \vec{x}^{(t)} - [\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]. \quad (10.38)$$

The formal algorithm, detailed in Algorithm 7, just repeats this iteration.

Algorithm 7 Newton's Method

```

1: function NEWTONMETHOD( $f, \vec{x}^{(0)}, T$ )
2:   for  $t = 0, 1, \dots, T - 1$  do
3:      $\vec{x}^{(t+1)} \leftarrow \vec{x}^{(t)} - [\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]$ 
4:   end for
5:   return  $\vec{x}^{(T)}$ 
6: end function

```

We call the vector $[\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]$ the Newton direction. Here, we do not choose a step-size η . Instead, we take a *full* step in the Newton direction towards the minimizer of the quadratic approximation of the objective function. This is the basic version of Newton's method; it is *not* guaranteed to converge in general. To achieve convergence, we can introduce a step-size $\eta > 0$ to the Newton update, obtaining the so-called *damped* Newton's method, which has the iteration

$$\vec{x}^{(t+1)} = \vec{x}^{(t)} - \eta [\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]. \quad (10.39)$$

A full algorithm is very similar to Algorithm 7 and is provided in Algorithm 8.

Algorithm 8 Damped Newton's Method

```

1: function DAMPEDNEWTONMETHOD( $f, \vec{x}^{(0)}, \eta, T$ )
2:   for  $t = 0, 1, \dots, T - 1$  do
3:      $\vec{x}^{(t+1)} \leftarrow \vec{x}^{(t)} - \eta [\nabla^2 f(\vec{x}^{(t)})]^{-1} [\nabla f(\vec{x}^{(t)})]$ 
4:   end for
5:   return  $\vec{x}^{(T)}$ 
6: end function

```

We will not discuss convergence proofs of Newton's method in this course; you may use [7] for further reading.

All of our discussion thus far has been under the assumption $\nabla^2 f(\vec{x}^{(t)}) \succ 0$. If the Hessian is not positive definite, one may adapt the algorithm accordingly, forming a new class of methods called *quasi-Newton's methods*. Discussion of quasi-Newton's methods are out of scope of the course.

Finally, we discuss the algorithmic complexity of Newton's method. In every iteration, we need to compute and invert the Hessian $\nabla^2 f(\vec{x}^{(t)})$ to obtain the search direction. This is much more expensive than computing the gradient $\nabla f(\vec{x}^{(t)})$, which is used both in the gradient descent method and in Newton's method. However, this expensive step is not without justification; in many convex optimization problems, Newton's method can be shown to converge to the optimal solution much faster (i.e., in fewer iterations) than gradient descent.

10.3 Newton's Method with Linear Equality Constraints

Our derivation of Newton's method can be used to handle equality-constrained optimization problems. Let $A \in \mathbb{R}^{m \times n}$ and $\vec{y} \in \mathbb{R}^m$, and let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a twice-differentiable strictly convex function with positive definite Hessian. Consider the following convex optimization problem:

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}). \quad (10.40)$$

$$\text{s.t. } A\vec{x} = \vec{y}. \quad (10.41)$$

We will use the same approach as with the unconstrained Newton's method, that is, we will take the second-order Taylor approximation around $\vec{x}^{(t)}$ and minimize it over the constraint set $\Omega \doteq \{\vec{x} \in \mathbb{R}^n : A\vec{x} = \vec{y}\}$. This method gives the following constrained convex quadratic program:

$$\min_{\vec{x} \in \mathbb{R}^n} f(\vec{x}^{(t)}) + [\nabla f(\vec{x}^{(t)})]^\top (\vec{x} - \vec{x}^{(t)}) + \frac{1}{2} (\vec{x} - \vec{x}^{(t)})^\top [\nabla^2 f(\vec{x}^{(t)})] (\vec{x} - \vec{x}^{(t)}) \quad (10.42)$$

$$\text{s.t. } A\vec{x} = \vec{y}. \quad (10.43)$$

Note that the quadratic program is convex and, if the original problem is feasible (i.e., Ω is nonempty) that strong duality holds by Slater's condition. Thus, we can solve this QP by solving the KKT conditions, as they are necessary and sufficient for global optimality. We begin by writing the Lagrangian $L: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ associated with this quadratic program, which is defined as

$$L(\vec{x}, \vec{\nu}) = f(\vec{x}^{(t)}) + [\nabla f(\vec{x}^{(t)})]^\top (\vec{x} - \vec{x}^{(t)}) + \frac{1}{2} (\vec{x} - \vec{x}^{(t)})^\top [\nabla^2 f(\vec{x}^{(t)})] (\vec{x} - \vec{x}^{(t)}) + \vec{\nu}^\top (A\vec{x} - \vec{y}). \quad (10.44)$$

Suppose that $(\vec{x}^*, \vec{\nu}^*)$ are globally optimal for the constrained quadratic program. Then they must satisfy the KKT conditions, which are:

- Primal feasibility:

$$A\vec{x}^* = \vec{y}. \quad (10.45)$$

- Stationarity/first-order condition:

$$\vec{0} = \nabla_{\vec{x}} L(\vec{x}^*, \vec{v}^*) \quad (10.46)$$

$$= \nabla f(\vec{x}^{(t)}) + [\nabla^2 f(\vec{x}^{(t)})](\vec{x}^* - \vec{x}^{(t)}) + A^\top \vec{v}. \quad (10.47)$$

Let us define a vector $\vec{v}^{(t)} = \vec{x}^* - \vec{x}^{(t)}$. Since $\vec{x}^{(t)}$ is feasible, we have $A\vec{x}^{(t)} = \vec{y}$. Thus we have

$$A\vec{v}^{(t)} = A(\vec{x}^* - \vec{x}^{(t)}) \quad (10.48)$$

$$= A\vec{x}^* - A\vec{x}^{(t)} \quad (10.49)$$

$$= \vec{y} - \vec{y} \quad (10.50)$$

$$= \vec{0}. \quad (10.51)$$

Thus, if we write the system in terms of $\vec{v}^{(t)}$ instead of \vec{x}^* , we have the system of equations

$$\vec{0} = \nabla f(\vec{x}^{(t)}) + [\nabla^2 f(\vec{x}^{(t)})]\vec{v}^{(t)} + A^\top \vec{v}^* \quad (10.52)$$

$$\vec{0} = A\vec{v}^{(t)} \quad (10.53)$$

which can be expressed in matrix form as

$$\begin{bmatrix} \nabla^2 f(\vec{x}^{(t)}) & A^\top \\ A & 0 \end{bmatrix} \begin{bmatrix} \vec{v}^{(t)} \\ \vec{v} \end{bmatrix} = \begin{bmatrix} -\nabla f(\vec{x}^{(t)}) \\ 0 \end{bmatrix}. \quad (10.54)$$

After solving this system of equations for $\vec{v}^{(t)}$, our update rule becomes

$$\vec{x}^{(t+1)} = \vec{x}^* = \vec{x}^{(t)} + \vec{v}^{(t)}, \quad (10.55)$$

which is equivalent to setting the new iterate as the minimizer of the constrained QP. The formal iteration is given in Algorithm 9.

Algorithm 9 Newton's Method with Linear Equality Constraints

- 1: **function** EQUALITYCONSTRAINEDNEWTONMETHOD($f, \vec{x}^{(0)}, T, A, \vec{y}$)
 - 2: **for** $t = 0, 1, \dots, T - 1$ **do**
 - 3: Solve the system $\begin{bmatrix} \nabla^2 f(\vec{x}^{(t)}) & A^\top \\ A & 0 \end{bmatrix} \begin{bmatrix} \vec{v}^{(t)} \\ \vec{v} \end{bmatrix} = \begin{bmatrix} -\nabla f(\vec{x}^{(t)}) \\ 0 \end{bmatrix}$ for $\vec{v}^{(t)}$
 - 4: $\vec{x}^{(t+1)} \leftarrow \vec{x}^{(t)} + \vec{v}^{(t)}$
 - 5: **end for**
 - 6: **return** $\vec{x}^{(T)}$
 - 7: **end function**
-

There also exist damped versions of this algorithm, but their analysis is out of scope of the course.

10.4 (OPTIONAL) Interior Point Method

In the previous section we introduced Newton's method, which allows us to solve convex optimization problems with simple linear equality constraints. In this section, we will build on top of Newton's method, introducing a new class

of algorithms to handle convex optimization problems with inequality constraints. Precisely, we will introduce the *interior point method*, which allows us to solve convex optimization problems of the following form

$$p^* = \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) \quad (10.56)$$

$$\text{s.t. } f_i(\vec{x}) \leq 0, \quad \forall i = 1, 2, \dots, m \quad (10.57)$$

$$A\vec{x} = \vec{y}, \quad (10.58)$$

where f_0, f_1, \dots, f_m are all convex twice-differentiable functions. Interior point methods (IPM) are a class of algorithms which solves the problem (10.56) by solving a sequence of convex optimization problems with linear constraints using Newton's method. The key idea used in IPM is the *barrier function*, which we introduce next.

10.4.1 Barrier Functions

Consider the problem (10.56). Our goal is to eliminate the inequality constraints and convert the problem to an equality constrained problem to which Newton's method can be applied. One way to do so is to augment the inequality constraints to the objective using indicator functions, as we did when developing our theory of duality. More precisely, consider the following unconstrained problem, which we denote by \mathcal{P} :

$$\text{problem } \mathcal{P}: \quad p^* = \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) + \sum_{i=1}^m I(f_i(\vec{x})) \quad (10.59)$$

$$A\vec{x} = \vec{b}. \quad (10.60)$$

where $I: \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ is given by

$$I(z) = \mathbb{1}[z \leq 0] = \begin{cases} 0, & \text{if } z \leq 0 \\ +\infty, & \text{if } z > 0 \end{cases} \quad (10.61)$$

This gives us an optimization problem with only linear equality constraints that is *equivalent* to the original optimization problem (10.56) (i.e., they have the same solution). However, introducing the indicator function now makes the objective function non-differentiable so we can no longer apply Newton's method to solve this problem. To overcome this problem, we will instead approximate the indicator function with a differentiable function ϕ , which we call a *barrier function*.

There are several choices for ϕ , i.e., good approximations for the indicator function, but they must all have something in common. Namely, ϕ should be a convex increasing function on \mathbb{R}_{-} , such that $\lim_{z \nearrow 0} \phi(z) = +\infty$, just like the indicator function I . There are many candidate functions that satisfy these criteria. One of the most used barrier functions that we will introduce here is the logarithmic barrier function, which, for some $\alpha > 0$, takes the form

$$\phi_\alpha(z) = -\frac{1}{\alpha} \log(-z), \quad (10.62)$$

The parameter α controls the accuracy of the approximation — as α grows larger, the logarithmic barrier function becomes a better and better approximation to the indicator function.

Using this logarithmic barrier, we can define an approximate optimization problem $\widehat{\mathcal{P}}(\alpha)$ to \mathcal{P} by the following:

$$\text{problem } \widehat{\mathcal{P}}(\alpha): \quad \widehat{p}_\alpha^* = \min_{\vec{x} \in \mathbb{R}^n} f_0(\vec{x}) + \sum_{i=1}^m \phi_\alpha(f_i(\vec{x})) \quad (10.63)$$

$$A\vec{x} = \vec{y}. \quad (10.64)$$

This optimization problem has a convex twice-differentiable objective function and linear equality constraints, so Newton's method can be applied.

10.4.2 Barrier Method

The problem $\widehat{\mathcal{P}}(\alpha)$ is just an approximation of the original problem \mathcal{P} . In particular, they do not necessarily have the same solution. Thus, the natural question to ask is how close the solutions of \mathcal{P} and $\widehat{\mathcal{P}}(\alpha)$ are. The choice of the parameter α will be crucial, since small α mean that ϕ_α does not behave like an indicator function and so $\widehat{\mathcal{P}}(\alpha)$ and \mathcal{P} are totally different in general, whereas large α offer much better approximations and thus allows us to expect much better solutions. Indeed, it is possible to show, with some additional assumptions [7], some more complicated bounds relating the distance of the solutions to \mathcal{P} to the solutions to $\widehat{\mathcal{P}}(\alpha)$, as a function of α . We do not go into this analysis here, but its main takeaway is that large values of α give us good approximate solutions.

One natural question is: why don't we just take the largest possible α and get a near-perfect approximation to the indicator functions, and thus the original program \mathcal{P} ? This actually may not be optimal from an algorithmic standpoint, as solving problem $\widehat{\mathcal{P}}(\alpha)$ using Newton's method becomes difficult for large values of α . This is because Newton's method relies on computing and inverting the Hessian of the objective, and with large values of α the Hessian changes rapidly for points near the boundary of the feasible set for the original problem, even points which may be the solution.

The barrier method overcomes this problem by solving the approximate problem for a relatively small value of α and obtain the approximate solution $\vec{x}^*(\alpha)$. The algorithm then refines this approximate solution by using it as an initial guess to solve the approximate problem with a larger value of α . This way, when the algorithm attempts to solve the difficult approximate problems $\widehat{\mathcal{P}}(\alpha)$ (as the value of α increases), it does so with a good initial guess. Newton's method converges extremely fast near the optimal solution, so this procedure greatly improves the convergence of Newton's method even when α is very large.¹

When solving constrained optimization problems, it is necessary for most algorithms to start with a feasible initialization $\vec{x}^{(0)}$. This is particularly important for the barrier method; in fact, it is necessary to start at a strictly feasible initial guess, so that the value of the approximate problem is finite and the derivative of the objective function is defined. In Algorithm 10 we present one instance of the barrier method in which the value of α is updated by scaling it with some constant $\mu > 1$.

Algorithm 10 Barrier Method.

```

1: function BARRIERMETHOD( $f_0, f_1, \dots, f_m, \vec{x}^{(0)}, \alpha^{(0)}, A, \vec{y}, \mu, T$ )
2:   for  $t = 1, \dots, T$  do
3:      $\vec{x}^{(t)} \leftarrow$  Solution of  $\widehat{\mathcal{P}}(\alpha^{(t-1)})$  using Newton's method (Algorithm 9) starting at  $\vec{x}^{(t-1)}$ 
4:      $\alpha^{(t)} \leftarrow \mu \alpha^{(t-1)}$ 
5:   end for
6:   return  $\vec{x}^{(T)}$ 
7: end function

```

¹This easy-to-hard solution process is one instance of a more general algorithmic paradigm called *homotopy continuation* or *homotopy analysis*, which is used to precisely simulate very unstable dynamical systems.

Chapter 11

Applications

In this chapter, we will discuss some applications of the theory we have developed so far in this class. Our exploration will include deterministic control and the linear-quadratic regulator, stochastic control and the policy gradient algorithm, and support vector machines.

11.1 Deterministic Control and Linear-Quadratic Regulator

The first application we'll discuss is in the area of deterministic control.

Although control is usually thought of as related to motion, it can apply to any dynamical system where we can understand the state and its dependence on time based on some function such as $\vec{x}_{t+1} = \vec{f}(\vec{x}_t, \vec{u}_t)$ which takes in a state \vec{x}_t and a control input \vec{u}_t and outputs the next state \vec{x}_{t+1} . The goal of control is to choose the control inputs \vec{u}_t to achieve some desired behavior of the system.

Example 214 (Vertical Rocket System). For example, we can consider a vertical rocket. Our goal is to maximize its height by time T . Let $x_{1,t}$ denote its height, $x_{2,t}$ denote its vertical velocity, and $x_{3,t}$ denote the weight of the rocket (which we will approximate as the weight of the fuel), all at time t . The weight of the fuel will go down over time, and that can affect the rocket's velocity. The forces pushing the rocket down are drag and gravity, and the upward force comes from the rocket's thrust.

The forces at time t have the following expressions. (Here \dot{x} is the time-derivative of x .)

- Inertial force: $x_{3,t} \cdot \ddot{x}_{1,t} = x_{3,t} \cdot \dot{x}_{2,t}$.
- Drag: $c_D \rho(x_{1,t}) x_{2,t}^2$ where c_D is a numerical constant and $\rho(x_1)$ is the density of the air at height x_1 .
- Gravity: $g x_{3,t}$ where g is a numerical constant.
- Thrust: $c_H \dot{x}_{3,t}$ where c_H is a numerical constant.

Given the input $u_t = -\dot{x}_{3,t}$, we want to write our dynamical system in the standard form for continuous dynamics, i.e., $\dot{\vec{x}}_t = \vec{f}(\vec{x}_t, u_t)$. From the force expressions, we have

$$\dot{x}_{1,t} = x_{2,t} \tag{11.1}$$

$$\dot{x}_{2,t} = \frac{1}{x_{3,t}} [-c_D \rho(x_{1,t}) x_{2,t}^2 + c_H \dot{x}_{3,t} - g x_{3,t}] \tag{11.2}$$

$$\dot{x}_{3,t} = -u_t \tag{11.3}$$

Now we have a dynamical system of the form $\dot{\vec{x}}_t = \vec{f}(\vec{x}_t, u_t)$. Recall that we want to maximize the height of the rocket, $x_{1,T}$, at some terminal time $T > 0$, given an initial condition $\vec{x}_0 = \vec{\xi}$ for some $\vec{\xi} \in \mathbb{R}^3$. We can set up an optimization problem to determine the $(u_t)_{t \in [0, T]}$ which accomplishes this:

$$\begin{aligned} \max_{(u_t)_{t \in [0, T]}} \quad & x_{1,T} \\ \text{s.t.} \quad & \dot{\vec{x}}_t = \vec{f}(\vec{x}_t, u_t), \quad \forall t \in [0, T] \\ & \vec{x}_0 = \vec{\xi}. \end{aligned}$$

In practice, we can use a numerical solver to solve this problem; conceptually one can solve many such systems by hand using the so-called *calculus of variations*, which is a sort of infinite-dimensional optimization paradigm which we do not explore more in this course.

Even though our dynamics are continuous-time and complex, we can discretize and locally linearize our system, obtaining an approximate system which is *discrete linear time-invariant*, i.e., of the form

$$\vec{x}_{k+1} = A\vec{x}_k + B\vec{u}_k, \quad \forall k \in \{0, 1, \dots, K-1\}. \quad (11.4)$$

where A and B are matrices of the appropriate sizes. This is a linear system because it is linear in $(\vec{x}_k)_{k=0}^K$ and $(\vec{u}_k)_{k=0}^{K-1}$, which we conceptually think of as very long (but finite-length) vectors. It is time-invariant because the matrices A and B do not depend on the discrete-time index k .

This particular type of control problem is ubiquitous within science and engineering, and thus deserves a special name — it is called the *linear quadratic regulator* problem.

Let's formally define our linear quadratic regulator system.

Definition 215 (Linear Quadratic Regulator (LQR))

Let $K \geq 0$ be a positive integer. Let $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $Q, Q_f \in \mathbb{S}_+^n$, and $R \in \mathbb{S}_+^m$. Let $\vec{\xi} \in \mathbb{R}^n$. A *linear quadratic regulator* is an optimization problem of the form:

$$\begin{aligned} \min_{\substack{(\vec{x}_k)_{k=0}^K \in \mathbb{R}^{(K+1)n} \\ (\vec{u}_k)_{k=0}^{K-1} \in \mathbb{R}^{Km}}} \quad & \frac{1}{2} \sum_{k=0}^{K-1} (\vec{x}_k^\top Q \vec{x}_k + \vec{u}_k^\top R \vec{u}_k) + \frac{1}{2} \vec{x}_K^\top Q_f \vec{x}_K \\ \text{s.t.} \quad & \vec{x}_{k+1} = A\vec{x}_k + B\vec{u}_k, \quad \forall k \in \{0, \dots, K-1\} \\ & \vec{x}_0 = \vec{\xi}. \end{aligned} \quad (11.5)$$

Equation (11.5) is actually a quadratic program, since the objective is quadratic in the variables $(\vec{x}_k)_{k=0}^K$ and $(\vec{u}_k)_{k=0}^{K-1}$, and the constraints are linear. We are able to solve this with the methods we already know. However, this problem is very large, having $(K+1)n + Km$ variables (as well as $K+1$ constraints), and for n, m, K large this quickly becomes intractable. Our saving grace is that this problem has significant *additional structure*. The traditional way to solve it is using the dynamic programming approach and Bellman's equation. However, in this section, we will solve it using the KKT conditions and the Riccati equation.

Theorem 216 (Optimal Control in LQR is Linear)

An optimal control for the LQR problem (11.5) is linear in the state, i.e.,

$$\vec{u}_k^* = -R^{-1}B^\top (I + P_{k+1}BR^{-1}B^\top)^{-1}P_{k+1}A\vec{x}_k^*, \quad \forall k \in \{0, \dots, K-1\} \quad (11.6)$$

where P_k are given by the recurrence relation

$$P_K = Q_f \quad (11.7)$$

$$P_k = A^\top (I + P_{k+1} B R^{-1} B^\top)^{-1} P_{k+1} A + Q, \quad \forall k \in \{0, \dots, K-1\}. \quad (11.8)$$

Proof. The Lagrangian of Equation (11.5) is

$$L((\vec{x}_k)_{k=0}^K, (\vec{u}_k)_{k=0}^{K-1}, (\vec{\lambda}_k)_{k=1}^K, \vec{v}) = \frac{1}{2} \sum_{k=0}^{K-1} (\vec{x}_k^\top Q \vec{x}_k + \vec{u}_k^\top R \vec{u}_k) + \frac{1}{2} \vec{x}_K^\top Q_f \vec{x}_K \quad (11.9)$$

$$+ \sum_{k=0}^{K-1} \vec{\lambda}_{k+1}^\top (A \vec{x}_k + B \vec{u}_k - \vec{x}_{k+1}) + \vec{v}^\top (\vec{x}_0 - \vec{\xi}). \quad (11.10)$$

Since the objective function of Equation (11.5) is convex and the constraints are affine, Slater's condition holds automatically. Thus strong duality holds. Since Equation (11.5) is a convex problem and strong duality holds, the KKT conditions are necessary and sufficient for global optimality.

Let $((\vec{x}_k^*)_{k=0}^K, (\vec{u}_k^*)_{k=0}^{K-1}, (\vec{\lambda}_k^*)_{k=1}^K, \vec{v}^*)$ be globally optimal primal and dual variables for Equation (11.5), hence satisfying the KKT conditions. Then we have

1. Primal feasibility:

$$(a) \quad \vec{x}_{k+1}^* = A \vec{x}_k^* + B \vec{u}_k^* \text{ for all } k \in \{0, \dots, K-1\}.$$

$$(b) \quad \vec{x}_0^* = \vec{\xi}.$$

2. Dual feasibility: N/A since all constraints are equality constraints.

3. Complementary slackness: N/A since all constraints are equality constraints.

4. Lagrangian stationarity:

$$(a) \quad \vec{0} = \nabla_{\vec{x}_0} L((\vec{x}_k^*)_{k=0}^K, (\vec{u}_k^*)_{k=0}^{K-1}, (\vec{\lambda}_k^*)_{k=1}^K, \vec{v}^*) = Q \vec{x}_0^* + A^\top \vec{\lambda}_1^* + \vec{v}^*.$$

$$(b) \quad \vec{0} = \nabla_{\vec{x}_k} L((\vec{x}_k^*)_{k=0}^K, (\vec{u}_k^*)_{k=0}^{K-1}, (\vec{\lambda}_k^*)_{k=1}^K, \vec{v}^*) = Q \vec{x}_k^* + A^\top \vec{\lambda}_{k+1}^* - \vec{\lambda}_k^* \text{ for all } k \in \{1, \dots, K-1\}.$$

$$(c) \quad \vec{0} = \nabla_{\vec{x}_K} L((\vec{x}_k^*)_{k=0}^K, (\vec{u}_k^*)_{k=0}^{K-1}, (\vec{\lambda}_k^*)_{k=1}^K, \vec{v}^*) = Q_f \vec{x}_K^* - \vec{\lambda}_K^*.$$

$$(d) \quad \vec{0} = \nabla_{\vec{x}_k} L((\vec{x}_k^*)_{k=0}^K, (\vec{u}_k^*)_{k=0}^{K-1}, (\vec{\lambda}_k^*)_{k=1}^K, \vec{v}^*) = R \vec{u}_k^* + B^\top \vec{\lambda}_{k+1}^* \text{ for all } k \in \{0, \dots, K-1\}.$$

The stationarity condition gives an update dynamics for $\vec{\lambda}_k^*$, i.e.,

$$\vec{\lambda}_k^* = A^\top \vec{\lambda}_{k+1}^* + Q \vec{x}_k^*, \quad \forall k \in \{1, \dots, K-1\}, \quad (11.11)$$

$$\vec{\lambda}_K^* = Q_f \vec{x}_K^*. \quad (11.12)$$

However, this update dynamics goes *backwards* in time from $k = K$ to $k = 1$. This is in contrast to the update dynamics for \vec{x}_k which goes forwards in time.

When we find the $\vec{\lambda}_k^*$, we are able to find the \vec{u}_k^* , since a stationarity equation gives

$$\vec{u}_k^* = -R^{-1} B^\top \vec{\lambda}_{k+1}^*, \quad \forall k \in \{0, \dots, K-1\}. \quad (11.13)$$

This motivates solving for $\vec{\lambda}_k^*$, which we do via (backwards) induction. Our induction hypothesis is of the form

$$\vec{\lambda}_k^* = P_k \vec{x}_k^*, \quad (11.14)$$

which we aim to show for $k \in \{1, \dots, K\}$. The base case is $k = K$, whence we have $\vec{\lambda}_K^* = Q_f \vec{x}_K^*$, so that $P_K = Q_f$. For the inductive step, for $k \in \{1, \dots, K-1\}$ we have

$$\vec{\lambda}_{k+1}^* = P_{k+1} \vec{x}_{k+1}^* \quad (11.15)$$

$$= P_{k+1} (A \vec{x}_k^* + B \vec{u}_k^*) \quad (11.16)$$

$$= P_{k+1} (A \vec{x}_k^* - BR^{-1} B^\top \vec{\lambda}_{k+1}^*) \quad (11.17)$$

$$= P_{k+1} A \vec{x}_k^* - P_{k+1} BR^{-1} B^\top \vec{\lambda}_{k+1}^* \quad (11.18)$$

$$(I + P_{k+1} BR^{-1} B^\top) \vec{\lambda}_{k+1}^* = P_{k+1} A \vec{x}_k^* \quad (11.19)$$

$$\vec{\lambda}_{k+1}^* = (I + P_{k+1} BR^{-1} B^\top)^{-1} P_{k+1} A \vec{x}_k^*. \quad (11.20)$$

$$\vec{\lambda}_k^* = A^\top \vec{\lambda}_{k+1}^* + Q \vec{x}_k^* \quad (11.21)$$

$$= A^\top ((I + P_{k+1} BR^{-1} B^\top)^{-1} P_{k+1} A \vec{x}_k^*) + Q \vec{x}_k^* \quad (11.22)$$

$$= \underbrace{(A^\top (I + P_{k+1} BR^{-1} B^\top)^{-1} P_{k+1} A + Q)}_{\doteq P_k} \vec{x}_k^* \quad (11.23)$$

$$= P_k \vec{x}_k^*. \quad (11.24)$$

Above, to show that (11.20) follows from (11.19), we need to confirm that $I + P_{k+1} BR^{-1} B^\top$ is invertible. To show this, we will explicitly construct its inverse in terms of the inverse of other matrices which we know are invertible. First, observe that $R + B^\top P_{k+1} B$ is invertible, since it is symmetric positive definite: R is symmetric positive definite and $B^\top P_{k+1} B$ is symmetric positive semidefinite, so their sum is symmetric positive definite. Next, we claim that $(I + P_{k+1} BR^{-1} B^\top)^{-1} = I - P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top$. This follows from the [Sherman-Morrison-Woodbury identity](#), but for the sake of completeness we prove it here. Indeed,

$$(I + P_{k+1} BR^{-1} B^\top)(I - P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top) \quad (11.25)$$

$$= I + P_{k+1} BR^{-1} B^\top - (I + P_{k+1} BR^{-1} B^\top) P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top \quad (11.26)$$

$$= I + P_{k+1} BR^{-1} B^\top - P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top - P_{k+1} BR^{-1} B^\top P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top \quad (11.27)$$

$$= I + P_{k+1} BR^{-1} B^\top - P_{k+1} BR^{-1} R (R + B^\top P_{k+1} B)^{-1} B^\top - P_{k+1} BR^{-1} B^\top P_{k+1} B (R + B^\top P_{k+1} B)^{-1} B^\top \quad (11.28)$$

$$= I + P_{k+1} BR^{-1} B^\top - P_{k+1} BR^{-1} (R + B^\top P_{k+1} B) (R + B^\top P_{k+1} B)^{-1} B^\top \quad (11.29)$$

$$= I + P_{k+1} BR^{-1} B^\top - P_{k+1} BR^{-1} B^\top \quad (11.30)$$

$$= I. \quad (11.31)$$

This confirms that $I + P_{k+1} BR^{-1} B^\top$ is invertible. Thus we have for $k \in \{0, \dots, K-1\}$ that

$$\vec{u}_k^* = -R^{-1} B^\top \vec{\lambda}_{k+1}^* \quad (11.32)$$

$$= -R^{-1} B^\top (I + P_{k+1} BR^{-1} B^\top)^{-1} P_{k+1} A \vec{x}_k^* \quad (11.33)$$

as desired. \square

Now, note that we have written our recurrence for P_k in terms of P_{k+1} . Starting from P_K , we can compute each P_k in a backwards order, completely offline (i.e., without processing any iterations of the forward system). Once we have the P_k , we can then compute each \vec{x}_k^* , $\vec{\lambda}_k^*$, and \vec{u}_k^* directly. Therefore, we have a way to solve the LQR problem just by using matrix multiplication.

11.2 Support Vector Machines

Support vector machines are an application of optimization to the classical machine learning problem of binary classification. We will see that using optimization and a specific perspective induced by the KKT conditions will allow us to elegantly solve this classification problem.

Let's review the basics of classification. Suppose that we have some image, and we want to figure out what kind of animal it is. For example, we could be classifying it as a cat or a dog. Or, it could be X-ray images, and we want to figure out if there is a fracture present or not. Nowadays, we use a wide range of techniques such as deep neural networks to solve these tasks. These more advanced techniques build up from much more fundamental ideas such as support vector machines.

For the purpose of this class, we will focus on binary classification. In binary classification, we are given some data $\vec{x}_1, \dots, \vec{x}_n \in \mathbb{R}^d$ as well as their corresponding labels $y_1, \dots, y_n \in \{-1, +1\}$. We want to find a function $f: \mathbb{R}^d \rightarrow \{-1, 1\}$ such that $f(\vec{x}_i) = y_i$ for all i . We could try solving this problem using least squares, but in many cases the least squares solution is suboptimal at classification tasks, being too susceptible to outliers. Here, we will explore a different technique, called *support vector machines* (SVMs).

The basic premise of support-vector machines is that we want to find an *affine function* $g_{\vec{w}, b}: \vec{x} \mapsto \vec{w}^\top \vec{x} - b$ which (approximately) separates the data into their respective classes, i.e., $y_i = \text{sgn}(g_{\vec{w}, b}(\vec{x}_i))$ for all i , so $f = \text{sgn} \circ g_{\vec{w}, b}$. Geometrically, $g_{\vec{w}, b}$ can be thought of as a *separating hyperplane* for the two classes, in fact corresponding to the hyperplane $\mathcal{H}_{\vec{w}, b} \doteq \{\vec{x} \in \mathbb{R}^d \mid \vec{w}^\top \vec{x} = b\}$.

If $y_i = 1$ then we would like $g_{\vec{w}, b}(\vec{x}_i) > 0$. If $y_i = -1$ then we would like $g_{\vec{w}, b}(\vec{x}_i) < 0$. We can express these two desires as always wanting $y_i g_{\vec{w}, b}(\vec{x}_i) > 0$, combining these two cases into one inequality.

Hard-Margin SVM

First, let us work through the case that the data are *strictly linearly separable*; that is, there exists some \vec{w}, b such that $y_i g_{\vec{w}, b}(\vec{x}_i) > 0$ for all i . This hypothesis is unreasonable in many cases, but it allows us to build intuition for the problem. We will later remove this hypothesis, but many tools remain the same.

In this case, we would like to find one (\vec{w}, b) pair for which $y_i g_{\vec{w}, b}(\vec{x}_i) > 0$ for all i . However, there can be many such pairs. Thus, we have to determine which one we would like to pick.

One possible heuristic¹ is to pick a pair (\vec{w}, b) with the largest *margin* — that is, the distance between the hyperplane $\mathcal{H}_{\vec{w}, b}$ and the closest point towards it. Thus we want to solve the problem:

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} \min_{i \in \{1, \dots, n\}} \text{dist}(\mathcal{H}_{\vec{w}, b}, \vec{x}_i) \\ \text{s.t. } y_i g_{\vec{w}, b}(\vec{x}_i) > 0, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \tag{11.34}$$

This problem is called the *hard-margin SVM*, so named because we do not allow any misclassification of the training points — no training data can “cross the margin.”

Unfortunately, the problem in Equation (11.34) seems intractable. Let us go about simplifying it. First, the distance between point \vec{x} and $\mathcal{H}_{\vec{w}, b}$ is defined as

$$\text{dist}(\mathcal{H}_{\vec{w}, b}, \vec{x}) \doteq \frac{|\vec{w}^\top \vec{x} - b|}{\|\vec{w}\|_2} = \frac{|g_{\vec{w}, b}(\vec{x})|}{\|\vec{w}\|_2}. \tag{11.35}$$

¹This heuristic makes sense from the perspective of robustness and generalization. We haven't sampled all the data that exists, but we hypothesize that the data we haven't sampled is geometrically close to the data we have sampled. We want to make sure that the largest fraction possible of all data (sampled and unsampled) is correctly classified by the \vec{w} and b we learn. Thus, to capture the most unsampled data possible, we require the classification to be as robust to these geometric deviations as possible.

Thus² we obtain the slightly simplified, yet equivalent, problem

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} & \frac{1}{\|\vec{w}\|_2} \min_{i \in \{1, \dots, n\}} |g_{\vec{w}, b}(\vec{x}_i)| \\ \text{s.t.} & y_i g_{\vec{w}, b}(\vec{x}_i) > 0, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.36)$$

Adding the real-valued slack variable s to denote the minimizing component, we have

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R} \\ s \in \mathbb{R}_{++}}} & \frac{s}{\|\vec{w}\|_2} \\ \text{s.t.} & y_i g_{\vec{w}, b}(\vec{x}_i) > 0, \quad \forall i \in \{1, \dots, n\}, \\ & |g_{\vec{w}, b}(\vec{x}_i)| \geq s, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.37)$$

Since $y_i \in \{-1, +1\}$ and $s > 0$, we have that for *any* scalar u_i that $|u_i| \geq s$ and $y_i u_i > 0$ if and only if $y_i u_i \geq s$. This relation certainly holds for $u_i = g_{\vec{w}, b}(\vec{x}_i)$, and so the above problem simplifies to the following:

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R} \\ s \in \mathbb{R}_{++}}} & \frac{s}{\|\vec{w}\|_2} \\ \text{s.t.} & y_i g_{\vec{w}, b}(\vec{x}_i) \geq s, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.38)$$

By the form of $g_{\vec{w}, b}(\vec{x}) = \vec{w}^\top \vec{x} - b$, the above problem is equivalent to the following problem:

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R} \\ s \in \mathbb{R}_{++}}} & \frac{1}{\|\vec{w}\|_2 / s} \\ \text{s.t.} & y_i g_{\vec{w}/s, b/s}(\vec{x}_i) \geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.39)$$

Since the above problem only depends on the values of \vec{w}/s and b/s , where \vec{w}/s is still an optimization variable allowed to take arbitrary values in \mathbb{R}^d and b/s is still an optimization variable allowed to take arbitrary values in \mathbb{R} , we may as well remove s by setting it to 1, obtaining

$$\begin{aligned} \max_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} & \frac{1}{\|\vec{w}\|_2} \\ \text{s.t.} & y_i g_{\vec{w}, b}(\vec{x}_i) \geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.40)$$

Finally, we can obtain an equivalent problem which is a convex minimization problem by using the transformation $x \mapsto \frac{1}{2x^2}$, which is monotonically decreasing on \mathbb{R}_{++} , on the objective function, whence (after expanding the form of $g_{\vec{w}, b}$), we obtain the problem

$$\begin{aligned} \min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} & \frac{1}{2} \|\vec{w}\|_2^2 \\ \text{s.t.} & y_i (\vec{w}^\top \vec{x}_i - b) \geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.41)$$

The problem (11.41) is by far the most common and simplified form of the hard-margin SVM problem. Moreover, this is a *quadratic program* in (\vec{w}, b) , which we know how to solve.

²It may seem natural to make $g_{\vec{w}, b}$ more complex and thus obtain a recipe for learning more complicated types of classifiers, but note that the above simplification no longer holds if we do, so the final problem may be much more complex and not efficiently solvable.

Soft Margin SVM

Most of our real world data isn't actually linearly separable. In that case, our hard margin SVM problem would just be infeasible. But maybe we still want to try to separate the points, even if our classifier is not perfect on the training data.

To develop this relaxed SVM problem, let us consider the hard-margin SVM:

$$\begin{aligned} \min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} & \frac{1}{2} \|\vec{w}\|_2^2 \\ \text{s.t.} & y_i g_{\vec{w}, b}(\vec{x}_i) \geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.42)$$

It is a constrained optimization problem, so we can reformulate it into an unconstrained problem using indicator variables:

$$\min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} \left(\frac{1}{2} \|\vec{w}\|_2^2 + \sum_{i=1}^n \ell_{0-\infty}(1 - y_i g_{\vec{w}, b}(\vec{x}_i)) \right) \quad (11.43)$$

where the $\ell_{0-\infty}$ loss is defined by

$$\ell_{0-\infty}(z) = \begin{cases} 0, & z \leq 0 \\ \infty, & z > 0. \end{cases} \quad (11.44)$$

This corresponds to infinitely penalizing any violation of the margin. In order to relax the penalties for margin violation, we can instead consider finite penalties which increase as the degree of violation of the margin increase. Namely, we can consider constant multiples of the so-called *hinge* loss

$$\ell_{\text{hinge}}(z) = \begin{cases} 0, & z \leq 0 \\ z, & z > 0 \end{cases} = \max\{z, 0\}. \quad (11.45)$$

This relaxation yields the following program:

$$\min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} \left(\frac{1}{2} \|\vec{w}\|_2^2 + C \sum_{i=1}^n \ell_{\text{hinge}}(1 - y_i g_{\vec{w}, b}(\vec{x}_i)) \right), \quad (11.46)$$

or using the definition of the hinge loss, the equivalent program

$$\min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} \left(\frac{1}{2} \|\vec{w}\|_2^2 + C \sum_{i=1}^n \max\{1 - y_i g_{\vec{w}, b}(\vec{x}_i), 0\} \right). \quad (11.47)$$

We can introduce some slack variables $\vec{\xi}$ to model this maximization term and make the problem differentiable. After expanding the form of $g_{\vec{w}, b}$, we have the program

$$\min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R} \\ \vec{\xi} \in \mathbb{R}^n}} \frac{1}{2} \|\vec{w}\|_2^2 + C \sum_{i=1}^n \xi_i \quad (11.48)$$

$$\text{s.t.} \quad \xi_i \geq 0, \quad \forall i \in \{1, \dots, n\} \quad (11.49)$$

$$\xi_i \geq 1 - y_i (\vec{w}^\top \vec{x}_i - b), \quad \forall i \in \{1, \dots, n\}. \quad (11.50)$$

This is the usual form of the *soft margin SVM* problem, and the solutions are parameterized by C .

In accordance with our derivation, if C is large then we allow only small violations to the margin, because the second term becomes a better approximation of the sum of indicators in Equation (11.43). If C is small, then we allow larger violations to the margin.

Depending on the perspective, either the first term or the second term of the loss can be viewed as the regularizer. The first term works to maximize the margin, while the second term works to penalize the margin violations. Both work together to form an approximate maximum-margin classifier.

KKT Conditions

It turns out that we get significant insight into the solutions to the hard-margin and soft-margin SVM using the KKT conditions. First, let us consider the hard-margin SVM:

$$\begin{aligned} \min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R}}} & \frac{1}{2} \|\vec{w}\|_2^2 \\ \text{s.t.} & y_i(\vec{w}^\top \vec{x}_i - b) \geq 1, \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (11.51)$$

The Lagrangian is

$$L(\vec{w}, b, \vec{\lambda}) = \frac{1}{2} \|\vec{w}\|_2^2 + \sum_{i=1}^n \lambda_i (1 - y_i(\vec{w}^\top \vec{x}_i - b)). \quad (11.52)$$

This problem is convex, and the constraints are affine, so if the problem is feasible (i.e., the data are strictly linearly separable), then Slater's condition holds so strong duality holds. Since the problem is convex and strong duality holds, the KKT conditions are necessary and sufficient for optimality.

Suppose that $(\vec{w}^*, b^*, \vec{\lambda}^*)$ satisfy the KKT conditions. Then:

1. Primal feasibility: $y_i((\vec{w}^*)^\top \vec{x}_i - b^*) \geq 1$ for all $i \in \{1, \dots, n\}$.
2. Dual feasibility: $\lambda_i^* \geq 0$ for all $i \in \{1, \dots, n\}$.
3. Stationarity:

$$(a) \quad \vec{0} = \nabla_{\vec{w}} L(\vec{w}^*, b^*, \vec{\lambda}^*) = \vec{w}^* - \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i.$$

$$(b) \quad 0 = \nabla_b L(\vec{w}^*, b^*, \vec{\lambda}^*) = \sum_{i=1}^n \lambda_i^* y_i.$$

4. Complementary slackness: $\lambda_i^*(1 - y_i((\vec{w}^*)^\top \vec{x}_i - b^*)) = 0$ for all $i \in \{1, \dots, n\}$.

We say that (\vec{x}_i, y_i) is a *support vector* if $\lambda_i^* > 0$. To see why, we consider the following cases:

1. If $\lambda_i^* = 0$ then, since $\vec{w}^* = \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i$, we see that (\vec{x}_i, y_i) does not contribute to the optimal solution.
2. If $\lambda_i^* > 0$ then by complementary slackness we have $y_i((\vec{w}^*)^\top \vec{x}_i - b^*) = 1$. Thus \vec{x}_i is on the margin of the SVM. Furthermore, since $\vec{w}^* = \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i$, we see that (\vec{x}_i, y_i) does contribute to the optimal solution.

Now let us consider the analogous notion for soft-margin SVMs. Consider the soft-margin SVM problem:

$$\min_{\substack{\vec{w} \in \mathbb{R}^d \\ b \in \mathbb{R} \\ \vec{\xi} \in \mathbb{R}^n}} \frac{1}{2} \|\vec{w}\|_2^2 + C \sum_{i=1}^n \xi_i \quad (11.53)$$

$$\text{s.t.} \quad \xi_i \geq 0, \quad \forall i \in \{1, \dots, n\} \quad (11.54)$$

$$\xi_i \geq 1 - y_i(\vec{w}^\top \vec{x}_i - b), \quad \forall i \in \{1, \dots, n\}. \quad (11.55)$$

It has Lagrangian

$$L(\vec{w}, b, \vec{\xi}, \vec{\lambda}, \vec{\mu}) = \frac{1}{2} \|\vec{w}\|_2^2 + C \sum_{i=1}^n \xi_i + \sum_{i=1}^n \mu_i (-\xi_i) + \sum_{i=1}^n \lambda_i (1 - y_i(\vec{w}^\top \vec{x}_i - b) - \xi_i). \quad (11.56)$$

This problem is convex, and the constraints are affine; one can show that it is always feasible, so Slater's condition holds and strong duality holds. Since the problem is convex with strong duality, the KKT conditions are both necessary and sufficient for optimality.

Suppose that $(\vec{w}^*, b^*, \vec{\xi}^*, \vec{\lambda}^*, \vec{\mu}^*)$ satisfy the KKT conditions. Then:

1. Primal feasibility: $\xi_i^* \geq 0$ and $\xi_i^* \geq 1 - y_i((\vec{w}^*)^\top \vec{x}_i - b^*)$ for all $i \in \{1, \dots, n\}$.

2. Dual feasibility: $\lambda_i^* \geq 0$ and $\mu_i^* \geq 0$ for all $i \in \{1, \dots, n\}$.

3. Stationarity:

(a) $\vec{0} = \nabla_{\vec{w}} L(\vec{w}^*, b^*, \vec{\xi}^*, \vec{\lambda}^*, \vec{\mu}^*) = \vec{w}^* - \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i$.

(b) $0 = \nabla_b L(\vec{w}^*, b^*, \vec{\xi}^*, \vec{\lambda}^*, \vec{\mu}^*) = \sum_{i=1}^n \lambda_i^* y_i$.

(c) $\vec{0} = \nabla_{\vec{\xi}} L(\vec{w}^*, b^*, \vec{\xi}^*, \vec{\lambda}^*, \vec{\mu}^*) = C\vec{1} - \vec{\mu}^* - \vec{\lambda}^*$.

4. Complementary slackness: $\lambda_i^*(1 - y_i((\vec{w}^*)^\top \vec{x}_i - b) - \xi_i^*) = 0$ and $\mu_i^* \xi_i^* = 0$ for all $i \in \{1, \dots, n\}$.

Similarly to the case of the hard-margin SVM, we say that (\vec{x}_i, y_i) is a *support vector* if $\lambda_i > 0$. To see why, we consider the following cases:

1. If $\lambda_i^* = 0$ then $\mu_i^* = C$. Thus $\mu_i^* > 0$, so $\xi_i^* = 0$. Thus the point \vec{x}_i does not violate the margin. Also since $\vec{w}^* = \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i^*$, we see that (\vec{x}_i, y_i) does not contribute to the optimal solution.

2. If $\lambda_i^* = C$ then $\mu_i^* = 0$. Thus we cannot say anything about ξ_i^* . But since $\lambda_i^* > 0$, the other complementary slackness condition says that $y_i((\vec{w}^*)^\top \vec{x}_i - b) = 1 - \xi_i^*$. Thus (\vec{x}_i, y_i) is either on the margin or violates the margin. Also since $\vec{w}^* = \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i^*$, we see that (\vec{x}_i, y_i) does contribute to the optimal solution.

3. If $\lambda_i^* \in (0, C)$, then $\mu_i^* \in (0, C)$ as well. Thus by complementary slackness, we have $\xi_i^* = 0$. Applying this to the other complementary slackness condition, we have $y_i((\vec{w}^*)^\top \vec{x}_i - b) = 1$. Thus (\vec{x}_i, y_i) is exactly on the margin. Also since $\vec{w}^* = \sum_{i=1}^n \lambda_i^* y_i \vec{x}_i^*$, we see that (\vec{x}_i, y_i) does contribute to the optimal solution.

In general, the support vectors contribute to the optimal solution, and they are on/violate the margin.

Bibliography

- [1] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [2] G. Calafiore and L. El Ghaoui, *Optimization Models*. Cambridge University Press, 2014.
- [3] C. C. Pugh, *Real Mathematical Analysis*. Springer, 2002, vol. 2011.
- [4] P. Varaiya *et al.*, *Lecture notes on optimization*. Unpublished manuscript, University of California, Department of Electrical Engineering and Computer Science, 1998.
- [5] D. P. Bertsekas, “Nonlinear programming,” *Journal of the Operational Research Society*, vol. 48, no. 3, pp. 334–334, 1997.
- [6] G. Garrigos and R. M. Gower, “Handbook of convergence theorems for (stochastic) gradient methods,” *arXiv preprint arXiv:2301.11235*, 2023.
- [7] Y. Nesterov *et al.*, *Lectures on convex optimization*. Springer, 2018, vol. 137.