

Introduction to Applied Linear Algebra

Vectors, Matrices, and Least Squares

Stephen Boyd

*Department of Electrical Engineering
Stanford University*

Lieven Vandenberghe

*Department of Electrical and Computer Engineering
University of California, Los Angeles*

July 14, 2018



CAMBRIDGE
UNIVERSITY PRESS

For

Anna, Nicholas, and Nora

Daniël and Margriet

Contents

1	Vectors	1
	Exercises	2
2	Linear functions	9
	Exercises	10
3	Norm and distance	17
	Exercises	18
4	Clustering	33
	Exercises	34
5	Linear independence	37
	Exercises	38
6	Matrices	43
	Exercises	44
7	Matrix examples	53
	Exercises	54
8	Linear equations	63
	Exercises	64
9	Linear dynamical systems	75
	Exercises	76
10	Matrix multiplication	79
	Exercises	80
11	Matrix inverses	99
	Exercises	100
12	Least squares	113
	Exercises	114

13 Least squares data fitting	125
Exercises	126
14 Least squares classification	141
Exercises	142
15 Multi-objective least squares	151
Exercises	152
16 Constrained least squares	161
Exercises	162
17 Constrained least squares applications	173
Exercises	174
18 Nonlinear least squares	181
Exercises	182
19 Constrained nonlinear least squares	197
Exercises	198

Chapter 1

Vectors

Exercises

- 1.1 Vector equations.** Determine whether each of the equations below is true, false, or contains bad notation (and therefore does not make sense).

(a) $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} = (1, 2, 1).$

(b) $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} = [1, 2, 1].$

(c) $(1, (2, 1)) = ((1, 2), 1).$

Solution.

- (a) *This equation is valid notation and is true.* Both sides of the equation are 3-vectors with equal entries.
- (b) *This equation doesn't make sense.* The expression $[1, 2, 1]$ is not valid notation.
- (c) *This equation is valid notation and is true.* Both sides of the equation use stacked vector notation and are equivalent to the vector $(1, 2, 1)$.

- 1.2 Vector notation.** Which of the following expressions uses correct notation? When the expression does make sense, give its length. In the following, a and b are 10-vectors, and c is a 20-vector.

(a) $a + b - c_{3:12}.$

(b) $(a, b, c_{3:13}).$

(c) $2a + c.$

(d) $(a, 1) + (c_1, b).$

(e) $((a, b), a).$

(f) $[a \ b] + 4c.$

(g) $\begin{bmatrix} a \\ b \end{bmatrix} + 4c.$

Solution.

- (a) *Correct.* Each of the vectors has length 10, so you can add them.
- (b) *Correct.* The length is the sum of the lengths of parts that are stacked, $10 + 10 + 11 = 31$.
- (c) *Incorrect.* $2a$ has length 10, but c has length 20, so you cannot add them.
- (d) *Correct.* The size of both vectors is 11.
- (e) *Correct.* (a, b) is a stacked vector of size 20, and $((a, b), a)$ is a stacked vector of size 30.
- (f) *Incorrect.* $[a \ b]$ is not a vector (it is a 10×2 matrix, which we study later), while $4c$ is a 20-vector.
- (g) *Correct.* Stacking a and b gives a vector of length 20 which can be added to c .

- 1.3 Overloading.** Which of the following expressions uses correct notation? If the notation is correct, is it also unambiguous? Assume that a is a 10-vector and b is a 20-vector.

(a) $b = (0, a).$

(b) $a = (0, b).$

(c) $b = (0, a, 0)$.

(d) $a = 0 = b$.

Solution.

- (a) Correct and unambiguous. 0 is the zero vector of length 10.
- (b) Incorrect. The left-hand side has length 10. The length of the right-hand is at least 21.
- (c) Correct but ambiguous. It is impossible to determine the sizes of the zero vectors from the equation.
- (d) Incorrect. The first equality only makes sense if 0 is the zero vector of length 10. The second equality only makes sense if 0 is the zero vector of length 20.

1.4 Periodic energy usage. The 168-vector w gives the hourly electricity consumption of a manufacturing plant, starting on Sunday midnight to 1AM, over one week, in MWh (megawatt-hours). The consumption pattern is the same each day, *i.e.*, it is 24-periodic, which means that $w_{t+24} = w_t$ for $t = 1, \dots, 144$. Let d be the 24-vector that gives the energy consumption over one day, starting at midnight.

- (a) Use vector notation to express w in terms of d .
- (b) Use vector notation to express d in terms of w .

Solution.

- (a) $w = (d, d, d, d, d, d, d)$.
- (b) $d = w_{1:24}$. There are other solutions, for example, $d = w_{25:48}$.

1.5 Interpreting sparsity. Suppose the n -vector x is sparse, *i.e.*, has only a few nonzero entries. Give a short sentence or two explaining what this means in each of the following contexts.

- (a) x represents the daily cash flow of some business over n days.
- (b) x represents the annual dollar value purchases by a customer of n products or services.
- (c) x represents a portfolio, say, the dollar value holdings of n stocks.
- (d) x represents a bill of materials for a project, *i.e.*, the amounts of n materials needed.
- (e) x represents a monochrome image, *i.e.*, the brightness values of n pixels.
- (f) x is the daily rainfall in a location over one year.

Solution.

- (a) On most days the business neither receives nor makes cash payments.
- (b) The customer has purchased only a few of the n products.
- (c) The portfolio invests (long or short) in only a small number of stocks.
- (d) The project requires only a few types of materials.
- (e) Most pixels have a brightness value of zero, so the image consists of mostly black space. This is the case, for example, in astronomy, where many imaging methods leverage the assumption that there are a few light sources against a dark sky.
- (f) On most days it didn't rain.

1.6 Vector of differences. Suppose x is an n -vector. The associated vector of differences is the $(n-1)$ -vector d given by $d = (x_2 - x_1, x_3 - x_2, \dots, x_n - x_{n-1})$. Express d in terms of x using vector operations (*e.g.*, slicing notation, sum, difference, linear combinations, inner product). The difference vector has a simple interpretation when x represents a time series. For example, if x gives the daily value of some quantity, d gives the day-to-day changes in the quantity.

Solution. $d = x_{2:n} - x_{1:n-1}$.

- 1.7 Transforming between two encodings for Boolean vectors.** A Boolean n -vector is one for which all entries are either 0 or 1. Such vectors are used to encode whether each of n conditions holds, with $a_i = 1$ meaning that condition i holds. Another common encoding of the same information uses the two values -1 and $+1$ for the entries. For example the Boolean vector $(0, 1, 1, 0)$ would be written using this alternative encoding as $(-1, +1, +1, -1)$. Suppose that x is a Boolean vector with entries that are 0 or 1, and y is a vector encoding the same information using the values -1 and $+1$. Express y in terms of x using vector notation. Also, express x in terms of y using vector notation.

Solution. We have $y = 2x - \mathbf{1}$. To see this, we note that $y_i = 2x_i - 1$. When $x_i = 1$, we have $y_i = 2 \cdot 1 - 1 = +1$; when $x_i = 0$, we have $y_i = 2 \cdot 0 - 1 = -1$.

Conversely, we have $x = (1/2)(y + \mathbf{1})$.

- 1.8 Profit and sales vectors.** A company sells n different products or items. The n -vector p gives the profit, in dollars per unit, for each of the n items. (The entries of p are typically positive, but a few items might have negative entries. These items are called *loss leaders*, and are used to increase customer engagement in the hope that the customer will make other, profitable purchases.) The n -vector s gives the total sales of each of the items, over some period (such as a month), *i.e.*, s_i is the total number of units of item i sold. (These are also typically nonnegative, but negative entries can be used to reflect items that were purchased in a previous time period and returned in this one.) Express the total profit in terms of p and s using vector notation.

Solution. The profit for item i is $p_i s_i$, so the total profit is $\sum_i p_i s_i = p^T s$. In other words, the total profit is just the inner product of p and s .

- 1.9 Symptoms vector.** A 20-vector s records whether each of 20 different symptoms is present in a medical patient, with $s_i = 1$ meaning the patient has the symptom and $s_i = 0$ meaning she does not. Express the following using vector notation.

- (a) The total number of symptoms the patient has.
- (b) The patient exhibits five out of the first ten symptoms.

Solution.

- (a) The total number of symptoms the patient has is $\mathbf{1}^T s$.
- (b) The patient exhibits five out of the first ten symptoms can be expressed as $a^T s = 5$, where the 20-vector a is given by $a = (\mathbf{1}_{10}, 0_{10})$. (The subscripts give the dimensions.)

- 1.10 Total score from course record.** The record for each student in a class is given as a 10-vector r , where $r_1, \dots, r_8$ are the grades for the 8 homework assignments, each on a 0–10 scale, r_9 is the midterm exam grade on a 0–120 scale, and r_{10} is final exam score on a 0–160 scale. The student's total course score s , on a 0–100 scale, is based 25% on the homework, 35% on the midterm exam, and 40% on the final exam. Express s in the form $s = w^T r$. (That is, determine the 10-vector w .) You can give the coefficients of w to 4 digits after the decimal point.

Solution. To convert the total homework score to a scale of 0–100, we add up the raw scores and multiply by $100/80 = 1.25$, $1.25(r_1 + \dots + r_8)$. The midterm score on a 0–100 scale is $(100/120)r_9 = 0.8333r_9$, and the final exam score on a 0–100 scale is $(100/160)r_{10} = 0.625r_{10}$. We multiply these by the weights 25%, 35%, and 40% to get the total score,

$$\begin{aligned} s &= (0.25)(1.25)(r_1 + \dots + r_8) + (0.35)(0.8)r_9 + (0.45)(0.625)r_{10} \\ &= (0.3125 \mathbf{1}_8, 0.2917, 0.25)^T r. \end{aligned}$$

So, $w = (0.3125 \mathbf{1}_8, 0.2917, 0.2500)$, where the subscript 8 on the ones vector tells us that it has size 8.

- 1.11 Word count and word count histogram vectors.** Suppose the n -vector w is the word count vector associated with a document and a dictionary of n words. For simplicity we will assume that all words in the document appear in the dictionary.

- (a) What is $\mathbf{1}^T w$?
- (b) What does $w_{282} = 0$ mean?
- (c) Let h be the n -vector that gives the histogram of the word counts, *i.e.*, h_i is the fraction of the words in the document that are word i . Use vector notation to express h in terms of w . (You can assume that the document contains at least one word.)

Solution.

- (a) $\mathbf{1}^T w$ is total number of words in the document.
- (b) $w_{282} = 0$ means that word 282 does not appear in the document.
- (c) $h = w/(\mathbf{1}^T w)$. We simply divide the word count vector by the total number of words in the document.

- 1.12 Total cash value.** An international company holds cash in five currencies: USD (US dollar), RMB (Chinese yuan), EUR (euro), GBP (British pound), and JPY (Japanese yen), in amounts given by the 5-vector c . For example, c_2 gives the number of RMB held. Negative entries in c represent liabilities or amounts owed. Express the total (net) value of the cash in USD, using vector notation. Be sure to give the size and define the entries of any vectors that you introduce in your solution. Your solution can refer to currency exchange rates.

Solution. The total value held in currency i , in USD, is given by $a_i c_i$, where a_i is the USD value of one unit of currency i . (So $a_i = 1$.) The total USD value is then $a_1 c_1 + \cdots + a_5 c_5 = a^T c$, the inner product of the exchange rate vector a and the currency holdings vector c .

- 1.13 Average age in a population.** Suppose the 100-vector x represents the distribution of ages in some population of people, with x_i being the number of $i-1$ year olds, for $i = 1, \dots, 100$. (You can assume that $x \neq 0$, and that there is no one in the population over age 99.) Find expressions, using vector notation, for the following quantities.

- (a) The total number of people in the population.
- (b) The total number of people in the population age 65 and over.
- (c) The average age of the population. (You can use ordinary division of numbers in your expression.)

Solution.

- (a) The total population is $\mathbf{1}^T x$.
- (b) The total number of people aged 65 or over is given by $a^T x$, where $a = (0_{65}, \mathbf{1}_{35})$. (The subscripts give the dimensions of the zero and ones vectors.)
- (c) The sum of the ages across the population is $(0, 1, 2, \dots, 99)^T x$. And so the average age is given by

$$\frac{(0, 1, 2, \dots, 99)^T x}{\mathbf{1}^T x}.$$

- 1.14 Industry or sector exposure.** Consider a set of n assets or stocks that we invest in. Let f be an n -vector that encodes whether each asset is in some specific industry or sector, *e.g.*, pharmaceuticals or consumer electronics. Specifically, we take $f_i = 1$ if asset i is in the sector, and $f_i = 0$ if it is not. Let the n -vector h denote a portfolio, with h_i the dollar value held in asset i (with negative meaning a short position). The inner product $f^T h$ is called the (dollar value) *exposure* of our portfolio to the sector. It gives the net dollar value of the portfolio that is invested in assets from the sector. A portfolio h is called *neutral* (to a sector or industry) if $f^T h = 0$.

A portfolio h is called *long only* if each entry is nonnegative, *i.e.*, $h_i \geq 0$ for each i . This means the portfolio does not include any short positions.

What does it mean if a long-only portfolio is neutral to a sector, say, pharmaceuticals? Your answer should be in simple English, but you should back up your conclusion with an argument.

Solution. If h is neutral to a sector represented by f , we have $f^T h = 0$. But this means

$$f_1 h_1 + \cdots + f_n h_n = 0.$$

Each term in this sum is the product of two nonnegative numbers, and is nonnegative. It follows that each term must be zero, *i.e.*, $f_i h_i = 0$ for $i = 1, \dots, n$. This in turn means that when asset i is in the sector, so $f_i = 1$, we must have $h_i = 0$. In other words: A long only portfolio is neutral to a sector only if it does not invest in any assets in that sector.

- 1.15 Cheapest supplier.** You must buy n raw materials in quantities given by the n -vector q , where q_i is the amount of raw material i that you must buy. A set of K potential suppliers offer the raw materials at prices given by the n -vectors $p_1, \dots, p_K$. (Note that p_k is an n -vector; $(p_k)_i$ is the price that supplier k charges per unit of raw material i .) We will assume that all quantities and prices are positive.

If you must choose just one supplier, how would you do it? Your answer should use vector notation.

A (highly paid) consultant tells you that you might do better (*i.e.*, get a better total cost) by splitting your order into two, by choosing two suppliers and ordering $(1/2)q$ (*i.e.*, half the quantities) from each of the two. He argues that having a diversity of suppliers is better. Is he right? If so, explain how to find the two suppliers you would use to fill half the order.

Solution. If we place the order with supplier i , the total price of the order is given by the inner product $p_i^T q$. We find the cheapest supplier by calculating $p_i^T q$ for $i = 1, \dots, K$, and finding the minimum of these K numbers.

Suppose the cheapest and second cheapest supplier are suppliers i and j . By splitting the order in two, the total price is $(p_i^T q + p_j^T q)/2$. This is never less than when we place the order with one supplier. However, ordering from more than one supplier can be advantageous for other reasons than cost.

- 1.16 Inner product of nonnegative vectors.** A vector is called *nonnegative* if all its entries are nonnegative.

- Explain why the inner product of two nonnegative vectors is nonnegative.
- Suppose the inner product of two nonnegative vectors is zero. What can you say about them? Your answer should be in terms of their respective sparsity patterns, *i.e.*, which entries are zero and nonzero.

Solution.

- Let x and y be nonnegative n -vectors. The inner product

$$x^T y = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n$$

is nonnegative because each term in the sum is nonnegative.

- For each k , $x_k = 0$ or $y_k = 0$ (or both). Therefore only the following three combinations of zero-positive patterns are positive (here $+$ stands for a positive entry):

x_k	y_k
+	0
0	+
0	0

- 1.17 Linear combinations of cash flows.** We consider cash flow vectors over T time periods, with a positive entry meaning a payment received, and negative meaning a payment made. A (unit) *single period loan*, at time period t , is the T -vector l_t that corresponds

to a payment received of \$1 in period t and a payment made of $\$(1+r)$ in period $t+1$, with all other payments zero. Here $r > 0$ is the interest rate (over one period).

Let c be a $\$1$ $T-1$ period loan, starting at period 1. This means that \$1 is received in period 1, $\$(1+r)^{T-1}$ is paid in period T , and all other payments (i.e., $c_2, \dots, c_{T-1}$) are zero. Express c as a linear combination of single period loans.

Solution. We are asked to write the T -vector

$$c = (1, 0, \dots, 0, -(1+r)^{T-1})$$

as a linear combination of the $T-1$ vectors

$$l_t = (0, \dots, 0, 1, -(1+r), 0, \dots, 0), \quad t = 1, \dots, T-1.$$

In the definition of l_t there are $t-1$ leading and $T-t-1$ trailing zeros, i.e., the element 1 is in position t . There is only one way to do this:

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ -(1+r)^{T-1} \end{bmatrix} = \begin{bmatrix} 1 \\ -(1+r) \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} + (1+r) \begin{bmatrix} 0 \\ 1 \\ -(1+r) \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} \\ + (1+r)^2 \begin{bmatrix} 0 \\ 0 \\ 1 \\ -(1+r) \\ \vdots \\ 0 \\ 0 \end{bmatrix} + \dots + (1+r)^{T-2} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 1 \\ -(1+r) \end{bmatrix}.$$

In other words,

$$c = l_1 + (1+r)l_2 + (1+r)^2l_3 + \dots + (1+r)^{T-2}l_{T-1}.$$

The coefficients in the linear combination are $1, 1+r, (1+r)^2, \dots, (1+r)^{T-1}$.

The idea is that you extend the length of an initial loan by taking out a new loan each period to cover the amount that you owe. So after taking out a loan for \$1 in period 1, you take out a loan for $\$(1+r)$ in period 2, and end up owing $\$(1+r)^2$ in period 3. Then you take out a loan for $\$(1+r)^2$ in period 3, and end up owing $\$(1+r)^3$ in period 4, and so on.

- 1.18 Linear combinations of linear combinations.** Suppose that each of the vectors $b_1, \dots, b_k$ is a linear combination of the vectors $a_1, \dots, a_m$, and c is a linear combination of $b_1, \dots, b_k$. Then c is a linear combination of $a_1, \dots, a_m$. Show this for the case with $m = k = 2$. (Showing it in general is not much more difficult, but the notation gets more complicated.)

Solution. We will consider the case $m = k = 2$. We have

$$b_1 = \beta_1 a_1 + \beta_2 a_2, \quad b_2 = \beta_3 a_1 + \beta_4 a_2.$$

Now assume that $c = \alpha_1 b_1 + \alpha_2 b_2$. Then we have

$$\begin{aligned} c &= \alpha_1 b_1 + \alpha_2 b_2 \\ &= \alpha_1(\beta_1 a_1 + \beta_2 a_2) + \alpha_2(\beta_3 a_1 + \beta_4 a_2) \\ &= (\alpha_1 \beta_1 + \alpha_2 \beta_3) a_1 + (\alpha_1 \beta_2 + \alpha_2 \beta_4) a_2. \end{aligned}$$

This shows that c is a linear combination of a_1 and a_2 .

- 1.19** *Auto-regressive model.* Suppose that $z_1, z_2, \dots$ is a time series, with the number z_t giving the value in period or time t . For example z_t could be the gross sales at a particular store on day t . An *auto-regressive* (AR) model is used to predict z_{t+1} from the previous M values, $z_t, z_{t-1}, \dots, z_{t-M+1}$:

$$\hat{z}_{t+1} = (z_t, z_{t-1}, \dots, z_{t-M+1})^T \beta, \quad t = M, M+1, \dots$$

Here $\hat{z}_{t+1}$ denotes the AR model's prediction of z_{t+1} , M is the memory length of the AR model, and the M -vector β is the AR model coefficient vector. For this problem we will assume that the time period is daily, and $M = 10$. Thus, the AR model predicts tomorrow's value, given the values over the last 10 days.

For each of the following cases, give a short interpretation or description of the AR model in English, without referring to mathematical concepts like vectors, inner product, and so on. You can use words like 'yesterday' or 'today'.

- (a) $\beta \approx e_1$.
- (b) $\beta \approx 2e_1 - e_2$.
- (c) $\beta \approx e_6$.
- (d) $\beta \approx 0.5e_1 + 0.5e_2$.

Solution.

- (a) The prediction is $\hat{z}_{t+1} \approx z_t$. This means that the prediction of tomorrow's value is, simply, today's value.
- (b) The prediction is $\hat{z}_{t+1} \approx z_t + (z_t - z_{t-1})$. In other words, the prediction is found by linearly extrapolating from yesterday's and today's values.
- (c) The prediction is $\hat{z}_{t+1} \approx z_{t-6}$, which is the value six days ago, which is the same as one week before tomorrow. For example, if today is Sunday we predict the value for Monday that is last Monday's value.
- (d) The prediction is $\hat{z}_{t+1} \approx 0.5z_t + 0.5z_{t-1}$, the average of today's and yesterday's values.

- 1.20** How many bytes does it take to store 100 vectors of length 10^5 ? How many flops does it take to form a linear combination of them (with 100 nonzero coefficients)? About how long would this take on a computer capable of carrying out 1 Gflop/s?

Solution.

- (a) $8 \times 100 \times 10^5 = 8 \cdot 10^7$ bytes.
- (b) $100 \times 10^5 + 99 \times 10^5 \approx 2 \cdot 10^7$ flops.
- (c) About 20 milliseconds.