

2.1 Introduction

Previously, we introduced vectors and matrices as a way of writing systems of linear equations more compactly, demonstrating through our tomography example that modeling a set of measurements as a system of equations can be a powerful tool. In this note, we are going to more thoroughly discuss how to perform computations with vectors and matrices so that in later notes we can consider properties of vectors and matrices and see how these can be used to help us understand real-world systems.

2.2 Introduction to Vectors


What is a vector? In the last note, we defined a vector as an ordered list of numbers. Generally speaking, if we are given a collection of n real numbers such as x_1, x_2, \dots, x_n , we can represent this collection as a single point in an n -dimensional real space, \mathbb{R}^n , denoted as a **vector** \vec{x} :

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

Each x_i (for i between 1 and n) is called a **component**, or **element**, of the vector. The elements of a vector do not have to be real — they could be complex-valued, meaning $\vec{x} \in \mathbb{C}^n$ — but the majority of vectors you will encounter in 16A will be real. The **size** of a vector is the number of components it contains. Two vectors \vec{x} and \vec{y} are said to be **equal**, $\vec{x} = \vec{y}$, if they have the same size and $x_i = y_i$ for all i .

Remark. Often, vectors are represented as letters in boldface (\mathbf{x}), or with a small arrow on top (\vec{x}). In these lecture notes we will use the latter arrow notation, and the size of the vector will be given (for example, $\vec{x} \in \mathbb{R}^3$ means that \vec{x} contains three real numbers).

Example 2.1 (3-D Vector): For $n = 3$, we could have $x_1 = -1$, $x_2 = 3.5$, and $x_3 = 1$, which can be represented as the vector $\vec{x} = \begin{bmatrix} -1 \\ 3.5 \\ 1 \end{bmatrix}$. This represents a point in 3-D space (\mathbb{R}^3), which can be drawn as an arrow from the origin to this point in space:



Vectors can be used to help write systems of equations compactly, but they are also useful for representing a multitude of things — anything that can be represented as an ordered list of numbers can be expressed as a vector. For instance, in the tomography example, we can write a vector to represent the amount of light absorbed by each bottle in a row or column.

2.2.1 Examples of Vectors

Example 2.2 (Position vs. time): $\vec{x} \in \mathbb{R}^n$ can represent samples of a quantity at n time points. Imagine a car moving along a line. Its position at time t_1, t_2, \dots, t_n can be represented with a vector:

$$\vec{x} = \begin{bmatrix} x_{t_1} \\ x_{t_2} \\ \vdots \\ x_{t_n} \end{bmatrix}$$

Here, x_{t_i} represents the position of the car at time t_i .

Example 2.3 (Quadrotor state): Vectors can be used to represent the *state* of a system, defined as follows:

Definition 2.1 (State): The minimum information you need to completely characterize a system at a given point in time, without any need for more information about the past of the system.

State is a powerful concept because it lets us separate the past from the future. The state completely captures the present — and the past can only affect the future through the present. We will revisit the concept of state in 16B, and you will see this in your homework.

As an example, consider modeling the dynamics of a quadrotor. The state of a quadrotor at a particular time can be summarized by its 3D position, angular position, velocity, and angular velocity, which can be represented as a vector $\vec{q} \in \mathbb{R}^{12}$, as illustrated in Figure 1.




Figure 1: A quadrotor (left) and the vector containing the information needed to determine its state (right). The 3D position of the quadrotor is captured in its x , y , and z coordinates, while its angular position is measured as its roll, pitch, and yaw. A dot above a variable implies taking a derivative with respect to time, giving the velocity of each component (for instance, $\dot{x} = \frac{dx}{dt}$ is the quadrotor's velocity in the x direction).

Example 2.4 (Vectors representing functions): Consider a simple game in which points are earned or lost by rolling a 6-sided die. If you roll a 5 or lower, you earn the number of points that you roll, but if you roll a

6, you earn -6 points. We could define a function that maps the die roll outcome, d , to the number of points earned by that roll, $p(d)$:

$$p(d) = \begin{cases} d & \text{if } d \leq 5 \\ -d & \text{if } d = 6 \end{cases}$$

We could also capture this information in a vector that summarizes all possible values of $p(d)$:

$$\begin{bmatrix} p(1) \\ p(2) \\ p(3) \\ p(4) \\ p(5) \\ p(6) \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ -6 \end{bmatrix}$$

Because the input domain of the function $p(d)$ is restricted to the integers 1 through 6, it is easy to see how it can be represented as a vector. It turns out that we can build a vector representation of any function with a numerical output — even functions that operate on an input domain of infinite size, such as all real numbers or all integers. Such functions can be represented by an *infinite-dimensional* vector. But for the most part, “vectors” in 16A will refer to finite-dimensional vectors.

Example 2.5 (Color): This is relevant to the lab we will do! The vector

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}, \quad \vec{x} \in \mathbb{R}^3$$

can represent a color, with its components giving red, green, and blue (RGB) intensity values. Typically, the entries of \vec{x} will fall in a range that reflects the range of the color (often 0 to 255)¹. However, the reality of color is quite subtle and involves psychophysics as well as ideas that are related to ideas of sampling that we will talk about in 16B.

Example 2.6 (Image): A grayscale image of $m \times n$ pixels is effectively a matrix with m rows and n columns, with entries corresponding to the grayscale levels at the pixel location. This matrix in $\mathbb{R}^{m \times n}$ could be represented by a vector of length mn (and therefore, a vector in \mathbb{R}^{mn}) by placing all pixel readings in a single list. This is explained in more detail in the figure below.

$$\begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1n} \\ x_{21} & x_{22} & \cdots & x_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mn} \end{bmatrix} \Rightarrow \begin{bmatrix} x_{11} \\ \vdots \\ x_{1n} \\ x_{21} \\ \vdots \\ x_{2n} \\ \vdots \\ x_{m1} \\ \vdots \\ x_{mn} \end{bmatrix}$$

What about a color image? The image could be stored in a vector in \mathbb{R}^{3mn} , because each of the mn pixels has values for red, green, and blue.

What about a color video? This could be stored in a vector in $\mathbb{R}^{3mn(\text{number of frames in the video})}$.

¹Why 0 to 255? We will cover this more in 16B, but generally speaking, the 256 integers in this range can be represented easily as an 8-bit binary number because $256 = 2^8$.

2.2.2 Special Vectors

Definition 2.2 (Zero Vector): A **zero vector** is a vector with all the components equal to zero, usually just represented as $\vec{0}$. You can usually tell the size of the zero vector $\vec{0}$ from the context: if $\vec{x} \in \mathbb{R}^n$ is added to $\vec{0}$, then $\vec{0}$ must also be in \mathbb{R}^n .

Definition 2.3 (Standard Unit Vector): A **standard unit vector** is a vector with all components equal to 0 except for one element, which is equal to 1. A standard unit vector where the i th position is equal to 1 is written as \vec{e}_i . We can denote the 3 standard unit vectors in \mathbb{R}^3 as:

$$\vec{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \vec{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \vec{e}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

When talking about standard unit vectors in the context of states, we might also use the word “pure” to refer to such states. This is because they only have one kind of component in them. Other states are mixtures of pure states.

2.3 Vector Operations

2.3.1 Vector Addition

Two vectors of the same size and in the same space (e.g. complex numbers, real numbers, etc.) can be added together by adding their corresponding components. For example, we can add two vectors in \mathbb{R}^3 :

$$\begin{bmatrix} -1 \\ 3.5 \\ 0 \end{bmatrix} + \begin{bmatrix} 2 \\ -1 \\ 3 \end{bmatrix} = \begin{bmatrix} -1+2 \\ 3.5-1 \\ 0+3 \end{bmatrix} = \begin{bmatrix} 1 \\ 2.5 \\ 3 \end{bmatrix}$$

Vector addition can be performed graphically as well. In \mathbb{R}^n , you place the tail of the first vector at the origin, and then place the tail of the second vector at the first vector’s head. The vector from the origin to the head of the second vector is the resulting vector sum. The following illustration shows how this can be done for two vectors, \vec{x} and \vec{y} , in \mathbb{R}^2 :




Figure 3: Adding two vectors in \mathbb{R}^2 . We can see that $\vec{x} + \vec{y} = \vec{y} + \vec{x}$.

Properties of Vector Addition: Many of the properties of addition you are already familiar with when adding individual numbers hold for vector addition as well. For three vectors $\vec{x}, \vec{y}, \vec{z} \in \mathbb{R}^n$ (and $\vec{0} \in \mathbb{R}^n$), the

following properties hold:

Commutativity (shown in Figure 3): $\vec{x} + \vec{y} = \vec{y} + \vec{x}$

Associativity: $(\vec{x} + \vec{y}) + \vec{z} = \vec{x} + (\vec{y} + \vec{z})$

Additive identity: $\vec{x} + \vec{0} = \vec{x}$

Additive inverse: $\vec{x} + (-\vec{x}) = \vec{0}$

2.3.2 Scalar Multiplication

We can multiply a vector by a number, called a **scalar**, that is in the same space as the vector (e.g. also in the reals, \mathbb{R} , or in the complex numbers, \mathbb{C}).

Definition 2.4 (Scalar): A **scalar** is a number. In mathematics and physics, scalars can be used to describe magnitude or used to scale things (e.g. cut every element of a vector in half by multiplying by 0.5, or flip the signs of each element in a vector by multiplying by -1).

To perform scalar multiplication, just multiply each of the components of the vector by the scalar:


$$-3 \times \begin{bmatrix} -1 \\ 3.5 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ -10.5 \\ 0 \end{bmatrix}$$

In general, for a scalar α and vector \vec{x} , this looks like

$$\alpha \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \vdots \\ \alpha x_n \end{bmatrix}$$

As an example, we can scale a vector $\vec{x} \in \mathbb{R}^2$ by 2 or -2:

$$\vec{x} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, 2\vec{x} = \begin{bmatrix} 2 \\ 2 \end{bmatrix}, -2\vec{x} = \begin{bmatrix} -2 \\ -2 \end{bmatrix}$$



Example 2.7 (Negative Vector): To obtain a negative vector $-\vec{x}$, where $\vec{x} \in \mathbb{R}^n$, we can just multiply the vector \vec{x} by the scalar -1 . In other words, $-\vec{x} = -1 \times \vec{x}$.

Example 2.8 (Zero Vector): We can obtain the zero vector by multiplying any vector by 0: $0\vec{x} = \vec{0}$

Properties of Scalar Multiplication: Just as multiplying a scalar by a scalar is associative ($(xy)z = x(yz)$), distributive ($x(y+z) = xy + xz$), and the multiplicative identity holds (multiplying any number by 1 returns

that original number), multiplying a scalar by a vector has the same properties:

$$\begin{aligned} \textbf{Associativity:} & \quad (\alpha\beta)\vec{x} = \alpha(\beta\vec{x}) \\ \textbf{Distributivity:} & \quad (\alpha + \beta)\vec{x} = \alpha\vec{x} + \beta\vec{x} \\ \textbf{Multiplicative identity:} & \quad 1\vec{x} = \vec{x} \end{aligned}$$

2.3.3 Vector Transpose

The transpose of a vector (or a matrix) is a crucial piece of notation for vector and matrix calculations. We

write the transpose of a vector \vec{x} as \vec{x}^T . If $\vec{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$, then $\vec{x}^T = [x_1 \ \cdots \ x_n]$. This is sometimes used so

that writing down long column vectors doesn't take up vertical space. The transpose of a column vector (a vertical vector) is called a row vector (a horizontal vector). The transpose of a row vector is called a column vector.

2.3.4 Vector-Vector Multiplication

While the operations we have discussed thus far are largely consistent with your intuition from working with scalars, vector-vector multiplication is not quite as straightforward. By convention, a row vector can only be multiplied by a column vector (and vice versa). Moreover, vector-vector multiplication is not commutative, so $\vec{y}^T \vec{x} \neq \vec{x} \vec{y}^T$. In fact, while $\vec{y}^T \vec{x}$ is only defined when both vectors have the same number of elements, $\vec{x} \vec{y}^T$ is defined for vectors of any size.

Multiplying a column vector by a row vector on the left is one way of computing the **inner product** or **dot product** of two vectors,² which we will discuss further in later notes. Below is an illustration of $\vec{y}^T \vec{x}$:

$$\vec{y}^T \vec{x} = [y_1 \ y_2 \ \cdots \ y_n] \times \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = y_1 x_1 + y_2 x_2 + \cdots + y_n x_n$$

Fundamentally, $\vec{y}^T \vec{x}$ is a sum of products — each element in \vec{y} is multiplied by the corresponding element in \vec{x} , and all of these products are added together. This form of vector-vector multiplication translates two vectors in \mathbb{R}^n to \mathbb{R} . Both vectors must contain n elements for their product to be defined.

Multiplying a row vector by a column vector on the left, on the other hand, generates a matrix. For instance, the product of $\vec{x} \in \mathbb{R}^n$ and $\vec{y}^T \in \mathbb{R}^m$ is given by:

$$\vec{x} \vec{y}^T = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} [y_1 \ y_2 \ \cdots \ y_m] = \begin{bmatrix} x_1 y_1 & x_1 y_2 & \cdots & x_1 y_m \\ x_2 y_1 & x_2 y_2 & \cdots & x_2 y_m \\ \vdots & \vdots & \vdots & \vdots \\ x_n y_1 & x_n y_2 & \cdots & x_n y_m \end{bmatrix}$$

The entry in row i and column j of $\vec{x} \vec{y}^T$ is the product of x_i and y_j . This product is perhaps best understood as a particular form of matrix-matrix multiplication, which we will discuss more in the next section.

²These terms are not identical mathematically, but will give the same result for the real vectors we will be working with in 16A.

2.4 Introduction to Matrices

We have provided a more extended look at vectors, and now we will take a more in-depth look at matrices. What is a matrix? Just as a vector as an ordered list of numbers, a matrix is an ordered list of vectors. Alternately, a matrix is a rectangular array of numbers — its elements are ordered in two dimensions. By convention, a matrix is said to be in $\mathbb{R}^{m \times n}$ if it has m rows and n columns. Here is an example of a matrix $A \in \mathbb{R}^{m \times n}$:

$$A = \begin{bmatrix} A_{11} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots \\ A_{m1} & \cdots & A_{mn} \end{bmatrix}$$

Remark. Matrices are often represented by capital letters (e.g. A), sometimes in boldface (e.g. \mathbf{A}). In these notes, we will use non-bolded capital letters to indicate matrices. Subscripts are typically used to specify an element of a matrix, with the first subscript corresponding the row index and the second corresponding to the column index.

A matrix is said to be **square** if $m = n$ (that is, the number of rows and number of columns are equal). Just as we could compute the **transpose** of a vector by transforming rows into columns, we can compute the transpose of a matrix, A^T , by making its rows into the columns of a new matrix (or vice versa):

$$A^T = \begin{bmatrix} A_{11} & \cdots & A_{1m} \\ \vdots & \ddots & \vdots \\ A_{n1} & \cdots & A_{nm} \end{bmatrix}$$

Mathematically, A^T is the $n \times m$ matrix given by $(A^T)_{ij} = A_{ji}$. A square matrix is said to be **symmetric** if $A = A^T$, which means that $A_{ij} = A_{ji}$ for all i and j .


2.4.1 Examples of Matrices

You have already seen some examples of matrices in previous notes, where we used matrices to represent the light absorbance of each bottle in our tomography example and as a shorthand to represent a system of linear equations when we discussed Gaussian elimination. There are many other ways matrices can be useful; we will provide a few examples here.

Example 2.9 (Illumination pattern): The imaging lab, in which individual pixels are illuminated by a projector, is very similar to the tomography example, in which individual bottles are illuminated depending on the angle of light being applied to the box. The pixel illumination pattern can similarly be represented as a square matrix of 1's and 0's, with a row for each photodiode measurement and a column for each pixel in the image. If a particular pixel is illuminated during a particular measurement, the matrix contains a 1 at the appropriate row and column index; otherwise, it contains a 0. In the tomography example, we set up a system of linear equations and built a matrix of coefficient weights corresponding to which juice bottle was included within a particular measurement.

Example 2.10 (Water pumps): Imagine a city with three water reservoirs, creatively labeled A , B , and C . At the end of each day, $\frac{1}{2}$ of the water in reservoir A is transferred to reservoir B while the remainder stays in A , $\frac{1}{4}$ of the water in B is transferred to A while the remainder goes to C , and the water in reservoir C is

split evenly between all three reservoirs. We can draw the water transfer as the following diagram, called a directed graphical model, where a directed arrow indicates that each evening, the labeled fraction of water moves from one reservoir to another:



We can describe this pump system as a matrix P with 3 rows and 3 columns, with entries that describe the fraction of water that flows from each reservoir (represented by columns) to another reservoir (represented by rows) each evening:

$$P = \begin{bmatrix} P_{A \rightarrow A} & P_{B \rightarrow A} & P_{C \rightarrow A} \\ P_{A \rightarrow B} & P_{B \rightarrow B} & P_{C \rightarrow B} \\ P_{A \rightarrow C} & P_{B \rightarrow C} & P_{C \rightarrow C} \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{3} \\ \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & \frac{3}{4} & \frac{1}{3} \end{bmatrix}$$

Later, we will see how this formulation of a matrix can be used to model the dynamics of systems.

2.4.2 Special Matrices

Definition 2.5 (Zero Matrix): Similar to a zero vector, a zero matrix is a matrix with all components equal to zero. It is usually just represented as 0, where its size is implied from context. Here is an example of a 2×2 zero matrix:

$$0 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

Definition 2.6 (Identity Matrix): An identity matrix, often written as I , is a square matrix whose diagonal elements are 1 and whose off-diagonal elements are all 0. Here is an example of the 3×3 identity matrix:

$$I = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Note that the column vectors (and the transpose of the row vectors) of an $n \times n$ identity matrix are the unit vectors in \mathbb{R}^n . The identity matrix is useful because multiplying a square matrix by an identity matrix of the same size will yield the same initial matrix: $AI = IA = A$ (we will discuss matrix-matrix multiplication further at the end of this note).

2.4.3 Column vs Row Perspective

The water reservoir matrix example makes it clear that both the rows and the columns of a matrix are interpretable. The rows represent the perspective of the water in each reservoir for the next day — each row tells what proportion of water a reservoir will draw from the other reservoirs. The columns represent the perspective of the current time’s water in a particular reservoir — each column indicates how that water will be distributed between the other reservoirs the next day.

This complements the interpretation of rows and columns that come from the system-of-linear-equations perspective of doing experiments. There, each row of the matrix represents a particular experiment that took one particular measurement. For a given row, the coefficient entries represent how much the corresponding state variable affects the outcome of this *particular* experiment. In contrast, the columns represent the influence of a particular state variable on the *collection* of experiments taken together. These perspectives come in handy in interpreting matrix multiplication.

2.5 Matrix Operations

2.5.1 Matrix Addition

Two matrices of the same size can be added together by adding their corresponding components. For instance, we can add two matrices A and B (both in $\mathbb{R}^{m \times n}$) as follows:

$$A + B = \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & & \vdots \\ A_{m1} & A_{m2} & \dots & A_{mn} \end{bmatrix} + \begin{bmatrix} B_{11} & B_{12} & \dots & B_{1n} \\ B_{21} & B_{22} & \dots & B_{2n} \\ \vdots & \vdots & & \vdots \\ B_{m1} & B_{m2} & \dots & B_{mn} \end{bmatrix} = \begin{bmatrix} A_{11} + B_{11} & A_{12} + B_{12} & \dots & A_{1n} + B_{1n} \\ A_{21} + B_{21} & A_{22} + B_{22} & \dots & A_{2n} + B_{2n} \\ \vdots & \vdots & & \vdots \\ A_{m1} + B_{m1} & A_{m2} + B_{m2} & \dots & A_{mn} + B_{mn} \end{bmatrix}$$

Note that if the two matrices are not the same size, there will not be a direct correspondence between matrix entries and we cannot add them. Here is a specific example of adding two 3×2 matrices:

$$\begin{bmatrix} -1 & 3 \\ 3.5 & 2 \\ 0 & -0.1 \end{bmatrix} + \begin{bmatrix} 2 & -1 \\ -1 & -2 \\ 3 & 0.1 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 2.5 & 0 \\ 3 & 0 \end{bmatrix}$$

Properties of Matrix Addition: Just as many of the properties of scalar addition hold for vector addition, many of the same properties hold for matrix addition as well. For three matrices A , B , and C of the same size (and a zero matrix of the same size), the following properties are true:

$$\begin{aligned} \text{Commutativity:} & \quad A + B = B + A \\ \text{Associativity:} & \quad (A + B) + C = A + (B + C) \\ \text{Additive identity:} & \quad A + 0 = A \\ \text{Additive inverse:} & \quad A + (-A) = 0 \end{aligned}$$

2.5.2 Scalar-Matrix Multiplication

As with scalar-vector multiplication, multiplying a matrix by a scalar requires multiplying each component of the matrix by that scalar. For instance, to multiply a scalar α by a matrix $A \in \mathbb{R}^{m \times n}$, we do the following:

$$\alpha A = \alpha \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & & \vdots \\ A_{m1} & A_{m2} & \dots & A_{mn} \end{bmatrix} = \begin{bmatrix} \alpha A_{11} & \alpha A_{12} & \dots & \alpha A_{1n} \\ \alpha A_{21} & \alpha A_{22} & \dots & \alpha A_{2n} \\ \vdots & \vdots & & \vdots \\ \alpha A_{m1} & \alpha A_{m2} & \dots & \alpha A_{mn} \end{bmatrix}$$

Here is a specific example for a 3×2 matrix:

$$(-3) \begin{bmatrix} -1 & 3 \\ 3.5 & 2 \\ 0 & -0.1 \end{bmatrix} = \begin{bmatrix} 3 & -9 \\ -10.5 & -6 \\ 0 & 0.3 \end{bmatrix}$$

Example 2.11 (Negative Matrix): To obtain a negative matrix $-A$, we can just multiply the original matrix A by the scalar -1 . In other words, $-A = -1 \times A$.

Example 2.12 (Zero Matrix): We can obtain the zero matrix by multiplying any matrix by 0: $0 \times A = 0$

Properties of Scalar-Matrix Multiplication: Scalar-matrix multiplication has many of the same properties as scalar-scalar multiplication and scalar-vector multiplication. Specifically, for scalars α, β and matrices $A, B \in \mathbb{R}^{m \times n}$, the following properties hold:

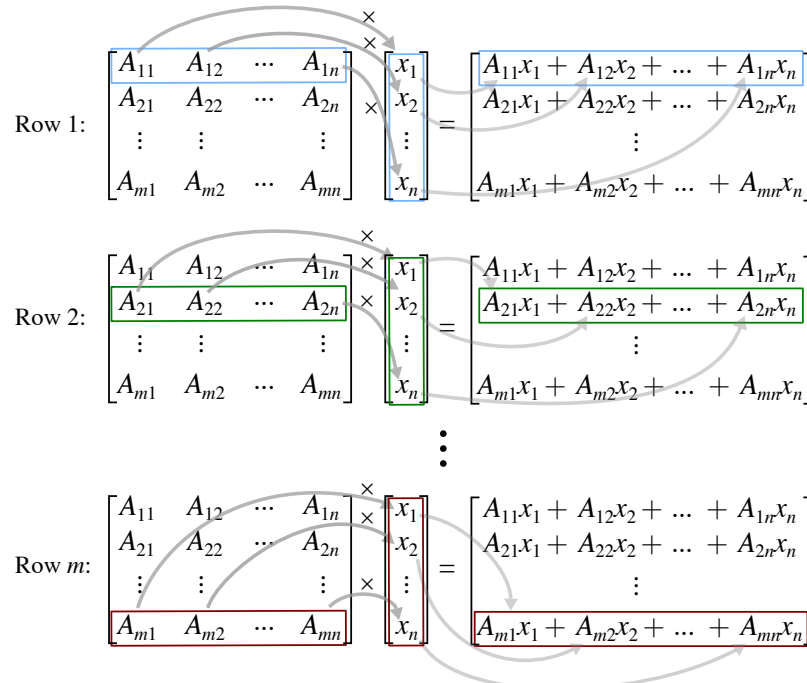
$$\begin{aligned} \text{Associativity:} \quad & (\alpha\beta) \times A = (\alpha) \times (\beta A) \\ \text{Distributivity:} \quad & (\alpha + \beta) \times A = \alpha A + \beta A \text{ and } \alpha(A + B) = \alpha A + \alpha B \\ \text{Multiplicative identity:} \quad & 1 \times A = A \end{aligned}$$

2.5.3 Matrix-Vector Multiplication

A matrix $A \in \mathbb{R}^{m \times n}$ and a vector $\vec{x} \in \mathbb{R}^n$ can be multiplied together to create a new vector of length m as follows:

$$A\vec{x} = \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & & \vdots \\ A_{m1} & A_{m2} & \dots & A_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} A_{11}x_1 + A_{12}x_2 + \dots + A_{1n}x_n \\ A_{21}x_1 + A_{22}x_2 + \dots + A_{2n}x_n \\ \vdots \\ A_{m1}x_1 + A_{m2}x_2 + \dots + A_{mn}x_n \end{bmatrix}$$

Each entry of the resulting vector is the inner product of the corresponding row of A with \vec{x} :



Mathematically speaking, if we say $A\vec{x} = \vec{b}$, with $A \in \mathbb{R}^{m \times n}$, $\vec{x} \in \mathbb{R}^n$, and $\vec{b} \in \mathbb{R}^m$, element i of \vec{b} can be calculated as follows:

$$b_i = \sum_{j=1}^n A_{ij}x_j$$

We can also see that the matrix-vector product can be rewritten as a sum of vectors, which can be expressed as a weighted sum of the columns of A :

$$\begin{aligned}
 A\vec{x} &= \begin{bmatrix} A_{11}x_1 + A_{12}x_2 + \cdots + A_{1n}x_n \\ A_{21}x_1 + A_{22}x_2 + \cdots + A_{2n}x_n \\ \vdots \\ A_{m1}x_1 + A_{m2}x_2 + \cdots + A_{mn}x_n \end{bmatrix} = \begin{bmatrix} A_{11}x_1 \\ A_{21}x_1 \\ \vdots \\ A_{m1}x_1 \end{bmatrix} + \begin{bmatrix} A_{12}x_2 \\ A_{22}x_2 \\ \vdots \\ A_{m2}x_2 \end{bmatrix} + \cdots + \begin{bmatrix} A_{1n}x_n \\ A_{2n}x_n \\ \vdots \\ A_{mn}x_n \end{bmatrix} \\
 &= x_1 \begin{bmatrix} A_{11} \\ A_{21} \\ \vdots \\ A_{m1} \end{bmatrix} + x_2 \begin{bmatrix} A_{12} \\ A_{22} \\ \vdots \\ A_{m2} \end{bmatrix} + \cdots + x_n \begin{bmatrix} A_{1n} \\ A_{2n} \\ \vdots \\ A_{mn} \end{bmatrix}
 \end{aligned}$$


How can this be interpreted? Consider the following example:

$$\begin{bmatrix} -1 & 3 \\ 3.5 & 2 \\ 0 & -0.1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ -0.1 \end{bmatrix}$$

Multiplying the standard unit vector by a matrix effectively selects a single column of the matrix — in this case, the second column. Matrix-vector multiplication is a way to compute a weighted sum of the columns in the matrix, where the weights are given by the vector that is being multiplied.

Now consider how matrix-vector multiplication applies to the examples you have seen. In the experiments and measurements perspective of matrices (the tomography example), we model the system of linear equations resulting from a set of measurements as $A\vec{x} = \vec{b}$. The columns of our matrix, A , represent the influence of each variable on the results of the experiment. The final measurements, \vec{b} , give a weighted combination of the influence of each variable, with weights given by the actual values that each variable has taken for a particular measurement. These weights are simply 1 or 0 in the tomography example, corresponding to whether or not a bottle is in the path of the light for a certain measurement.

Matrix-vector multiplication can also be employed in our water reservoir example. For instance, how would we calculate the next day's distribution of water in each reservoir? Below is the water flow graph for reference, where each blue circle represents one of three reservoirs (A , B , and C), and the labeled arrows indicate that a fraction of the water from one reservoir will be sent to another reservoir the next day:



Say reservoirs A , B , and C contain x_A , x_B , and x_C gallons of water, respectively. The amount of water each reservoir will contain the next day (x'_A , x'_B , and x'_C) can be calculated by adding together the fraction of water from each reservoir that flows into a given reservoir:

$$\begin{aligned}
 x'_A &= \frac{1}{2}x_A + \frac{1}{4}x_B + \frac{1}{3}x_C \\
 x'_B &= \frac{1}{2}x_A + \frac{1}{3}x_C \\
 x'_C &= \frac{3}{4}x_B + \frac{1}{3}x_C
 \end{aligned}$$

At this point, we can see that this is fundamentally a system of linear equations, which can be rewritten in the form $A\vec{x} = \vec{b}$:

$$\begin{bmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{3} \\ \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & \frac{3}{4} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} x_A \\ x_B \\ x_C \end{bmatrix} = \begin{bmatrix} x'_A \\ x'_B \\ x'_C \end{bmatrix}$$

From the equation above, we see that multiplying a vector of the initial contents of each reservoir by the pump transformation matrix will give a new vector with the distribution of water in each reservoir the next day. In other words, multiplying by the matrix tells us where the water ends up.

Here is where the weighted-column interpretation of matrix-vector multiplication comes into play — if the first, second, and third columns of the pump matrix indicate how water from reservoirs A , B , and C (respectively) are divided, then if all of the water starts out in a particular distribution between the three reservoirs, it should also end up in the appropriate combination of the three columns of the matrix.

2.5.4 Matrix-Matrix Multiplication (i.e., Transformation of Spaces)

Matrix-matrix multiplication is another powerful tool for modeling linear systems, which we will discuss further in later notes. As an example, two matrices A and B in $\mathbb{R}^{2 \times 2}$ can be multiplied as follows:

$$\begin{matrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} & \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} & = & \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix} \\ \text{A} & \text{B} & & \text{AB} \end{matrix}$$

Computationally, matrix-matrix multiplication involves taking the inner product of each row vector in A with each column vector in B , starting from the top row of matrix A and leftmost column of matrix B . Effectively, the left matrix is multiplied by each column vector in the second matrix to produce a new column of AB . Why columns and not rows? That's just convention. But this does lead to an important point about the dimensions of matrix-matrix multiplication.

To left-multiply a matrix B by another matrix A , the number of columns in A must equal the number of rows in B . Otherwise, the product $A \times B$ cannot be calculated. Moreover, if A is an $m \times n$ matrix and B is $n \times p$, the product $A \times B$ will have dimensions $m \times p$. A visual illustration of this can be seen here, where the left matrix is broken up into m row vectors and the right matrix is represented as p column vectors:

$$\begin{bmatrix} \text{---} \vec{r}_1^T \text{---} \\ \text{---} \vec{r}_2^T \text{---} \\ \vdots \\ \text{---} \vec{r}_m^T \text{---} \end{bmatrix} \times \begin{bmatrix} | & | & \dots & | \\ \vec{c}_1 & \vec{c}_2 & \dots & \vec{c}_p \\ | & | & \dots & | \end{bmatrix} = \begin{bmatrix} \vec{r}_1^T \vec{c}_1 & \vec{r}_1^T \vec{c}_2 & \dots & \vec{r}_1^T \vec{c}_p \\ \vec{r}_2^T \vec{c}_1 & \vec{r}_2^T \vec{c}_2 & \dots & \vec{r}_2^T \vec{c}_p \\ \vdots & \vdots & \dots & \vdots \\ \vec{r}_m^T \vec{c}_1 & \vec{r}_m^T \vec{c}_2 & \dots & \vec{r}_m^T \vec{c}_p \end{bmatrix}$$

In order for the inner product $\vec{r}_i^T \vec{c}_j$ to be defined, each row vector (\vec{r}_i^T) must have the same number of entries as each column vector (\vec{c}_j). As a result, matrix-matrix multiplication is typically not commutative — $A \times B$ does not necessarily equal $B \times A$. In fact, both quantities can only be calculated if the number of rows in A equals the number of columns in B **and** the number of rows in B equals the number of columns in A .

To conclude, we have a quick example of taking the product of two 2×2 matrices.

Example 2.13 (Matrix Multiplication):

$$\begin{bmatrix} 2 & 4 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} (2)(1) + (4)(3) & (2)(2) + (4)(4) \\ (3)(1) + (1)(3) & (3)(2) + (1)(4) \end{bmatrix} = \begin{bmatrix} 14 & 20 \\ 6 & 10 \end{bmatrix}$$

In a later note, we will further explore how matrix-matrix multiplication applies to linear transformations.