The linear algebra techniques you'll learn in this book are some of the most powerful mathematical modelling tools that exist. At the core of linear algebra lies a very simple idea: *linearity*. A function f is *linear* if it obeys the equation

$$f(a\mathbf{x}_1 + b\mathbf{x}_2) \;=\; af(\mathbf{x}_1) + bf(\mathbf{x}_2),$$

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are any two inputs of the function. We use the term *linear combination* to describe any expression constructed from a set of variables by multiplying each variable by a constant and adding the results. In the above equation, the linear combination $a\mathbf{x}_1 + b\mathbf{x}_2$ of the inputs $\mathbf{x}_1$ and $\mathbf{x}_2$ is transformed into the linear combination $af(\mathbf{x}_1) + bf(\mathbf{x}_2)$ of the outputs of the function $f(\mathbf{x}_1)$ and $f(\mathbf{x}_2)$. **Essentially, linear functions transform a linear combination of inputs into the *same* linear combination of outputs.** If the input to the linear function f consists of five parts $\mathbf{x}_1$ and three parts $\mathbf{x}_2$, then the output of the function will consist of five parts $f(\mathbf{x}_1)$ and three parts $f(\mathbf{x}_2)$. That's it, that's all! Now you know everything there is to know about linear algebra. The rest of the book is just details.

Linear models are super useful

A significant proportion of the math models used in science describe *linear relationships* between quantities. Mathematicians, scientists, engineers, and business analysts develop and use linear models to make sense of the systems they study. Linear models are popular because they are **easy to describe mathematically**. We can obtain the parameters of a linear model for a real-world system by analyzing the system's behaviour for relatively few inputs. Let's illustrate this important point with an example.

Example You're visiting an art gallery. Inside, the screen of a tablet computer is being projected onto a giant wall. Anything you draw on the tablet instantly appears projected onto the wall. However, the tablet's user interface doesn't give any indication about how to hold the tablet "right side up." How can you find the correct orientation of the tablet so your drawing won't appear rotated or upside-down?

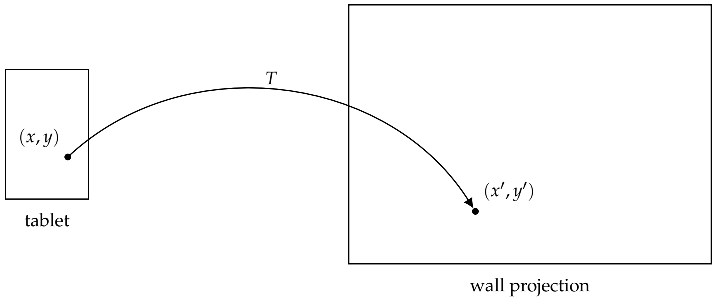

Figure 6: An unknown linear transformation T maps "tablet coordinates" to "wall coordinates." How can we characterize T?

The tablet's screen is a two-dimensional *input space* described by coordinates (x, y) and the wall projection is a two-dimensional *output space* described by wall coordinates (x', y'). You're looking for the unknown transformation T that maps the pixels of the tablet screen (the input space) to the projection on the wall (the output space):

$$(x, y) \xrightarrow{T} (x', y').$$

This task is directly analogous to the tasks scientists and engineers face every day when trying to model real-world systems by observing how systems transform inputs to outputs. If the unknown transformation T is linear, you can learn what it is very quickly, using only two swipes on the tablet screen.

To understand how T transforms screen coordinates (x, y) to wall coordinates (x', y'), you can use this two-step "probing" procedure:

1. Draw a horizontal line on the tablet to represent the x-direction in the input space $\rightarrow = (1, 0)$. You observe the output $\nearrow$ projected on the wall. This tells you horizontal lines are transformed to northeast diagonal lines in the wall-projection space.

2. Draw a vertical line in the y-direction $\uparrow = (0, 1)$ on the tablet. You observe the output $\nwarrow$ appears on the wall. This means vertical lines on the tablet screen turn into northwest diagonal lines when projected on the wall.

Here comes the interesting part: now that you know the outputs $\nearrow$ and $\nwarrow$ produced for the two input directions, you can **predict the linear transformation's output for any other input.** Let's look at the math equations that show why this is true.

Suppose you want to predict what will appear on the wall if you draw a line on the tablet in the direction $(3, 2)$. The coordinates $(3, 2)$ describe a swipe with length 3 units in the x-direction and 2 units

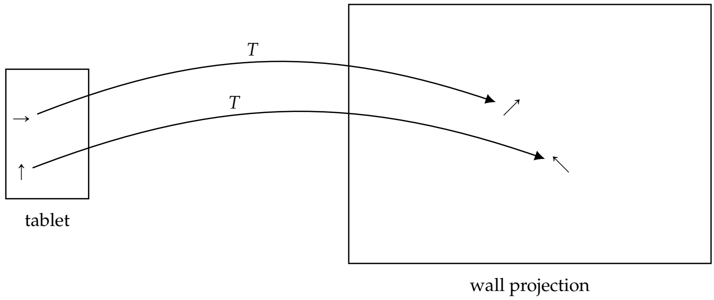

Figure 7: Drawing a short horizontal arrow → on the tablet screen results in a northeast diagonal projection on the wall ↗. Drawing a vertical arrow ↑ on the tablet results in a northwest diagonal line projected on the wall ↖.

in the y-direction. The input coordinates $(3, 2)$ can be written as $3(1, 0) + 2(0, 1) = 3 \rightarrow + 2 \uparrow$. Because you know T is linear, the wall projection of this input will have a length equal to 3 times the x-direction output ↗ plus 2 times the y-direction output ↖:

$$T\big(3 \rightarrow + 2 \uparrow\big) = 3T\big(\rightarrow\big) + 2T\big(\uparrow\big) = 3 \nearrow + 2 \nwarrow .$$

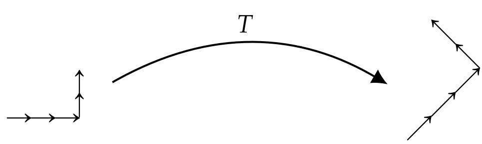

Figure 8: The linear transformation T maps the input $3 \rightarrow + 2 \uparrow$ to the output $3T(\rightarrow) + 2T(\uparrow) = 3 \nearrow + 2 \nwarrow$.

Knowing that the input → produces the output ↗ and the input ↑ produces the output ↖ allows you to determine the linear transformation's output for all other inputs. Every input (a, b) can be written as a linear combination: $(a, b) = a(1, 0) + b(0, 1) = a \rightarrow + b \uparrow$. Since you know T is linear, you know the corresponding output will be

$$T\big(a \rightarrow + b \uparrow\big) = aT\big(\rightarrow\big) + bT\big(\uparrow\big) = a \nearrow + b \nwarrow .$$

Since you can predict the output of T for all possible inputs, you have obtained a complete characterization of the linear transformation T.

The probing procedure we used to characterize the two-dimensional tablet-to-wall linear transformation (denoted $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$) can be used to study arbitrary linear transformations with n-dimensional inputs and m-dimensional outputs (denoted $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$). **Knowing the outputs of a linear transformation T for all "directions" in its input space gives us a complete characterization of T.**

TL;DR The linear property allows us to analyze multidimensional systems and processes by studying their effects on a small set of inputs. This is the essential reason linear models are used so widely in science. Without this linear structure, characterizing the behaviour of unknown input-output systems would be a much harder task.

Linear transformations

Linear transformations will be a central topic throughout this book. You can think of linear transformations as "vector functions" and understand their properties as analogous to the properties of the regular functions you're familiar with. The action of a function on a number is similar to the action of a linear transformation on a vector:

$$\begin{aligned}
\text{function } f : \mathbb{R} \to \mathbb{R} \; &\Leftrightarrow \; \text{linear transformation } T : \mathbb{R}^n \to \mathbb{R}^m \\
\text{input } x \in \mathbb{R} \; &\Leftrightarrow \; \text{input } \vec{x} \in \mathbb{R}^n \\
\text{output } f(x) \in \mathbb{R} \; &\Leftrightarrow \; \text{output } T(\vec{x}) \in \mathbb{R}^m \\
\text{inverse function } f^{-1} \; &\Leftrightarrow \; \text{inverse transformation } T^{-1} \\
\text{roots of } f \; &\Leftrightarrow \; \text{kernel of } T
\end{aligned}$$

Studying linear algebra will expose you to many new topics associated with linear transformations. You'll learn about concepts like vector spaces, projections, rotations, and orthogonalization procedures. Indeed, a first linear algebra course introduces many advanced, abstract ideas; yet all the new ideas you'll encounter can be seen as extensions of ideas you're already familiar with. Linear algebra is the vector-upgrade to your high school knowledge of functions.