kDimensions

A VISUAL INTRODUCTION TO

DEEP LEARNING

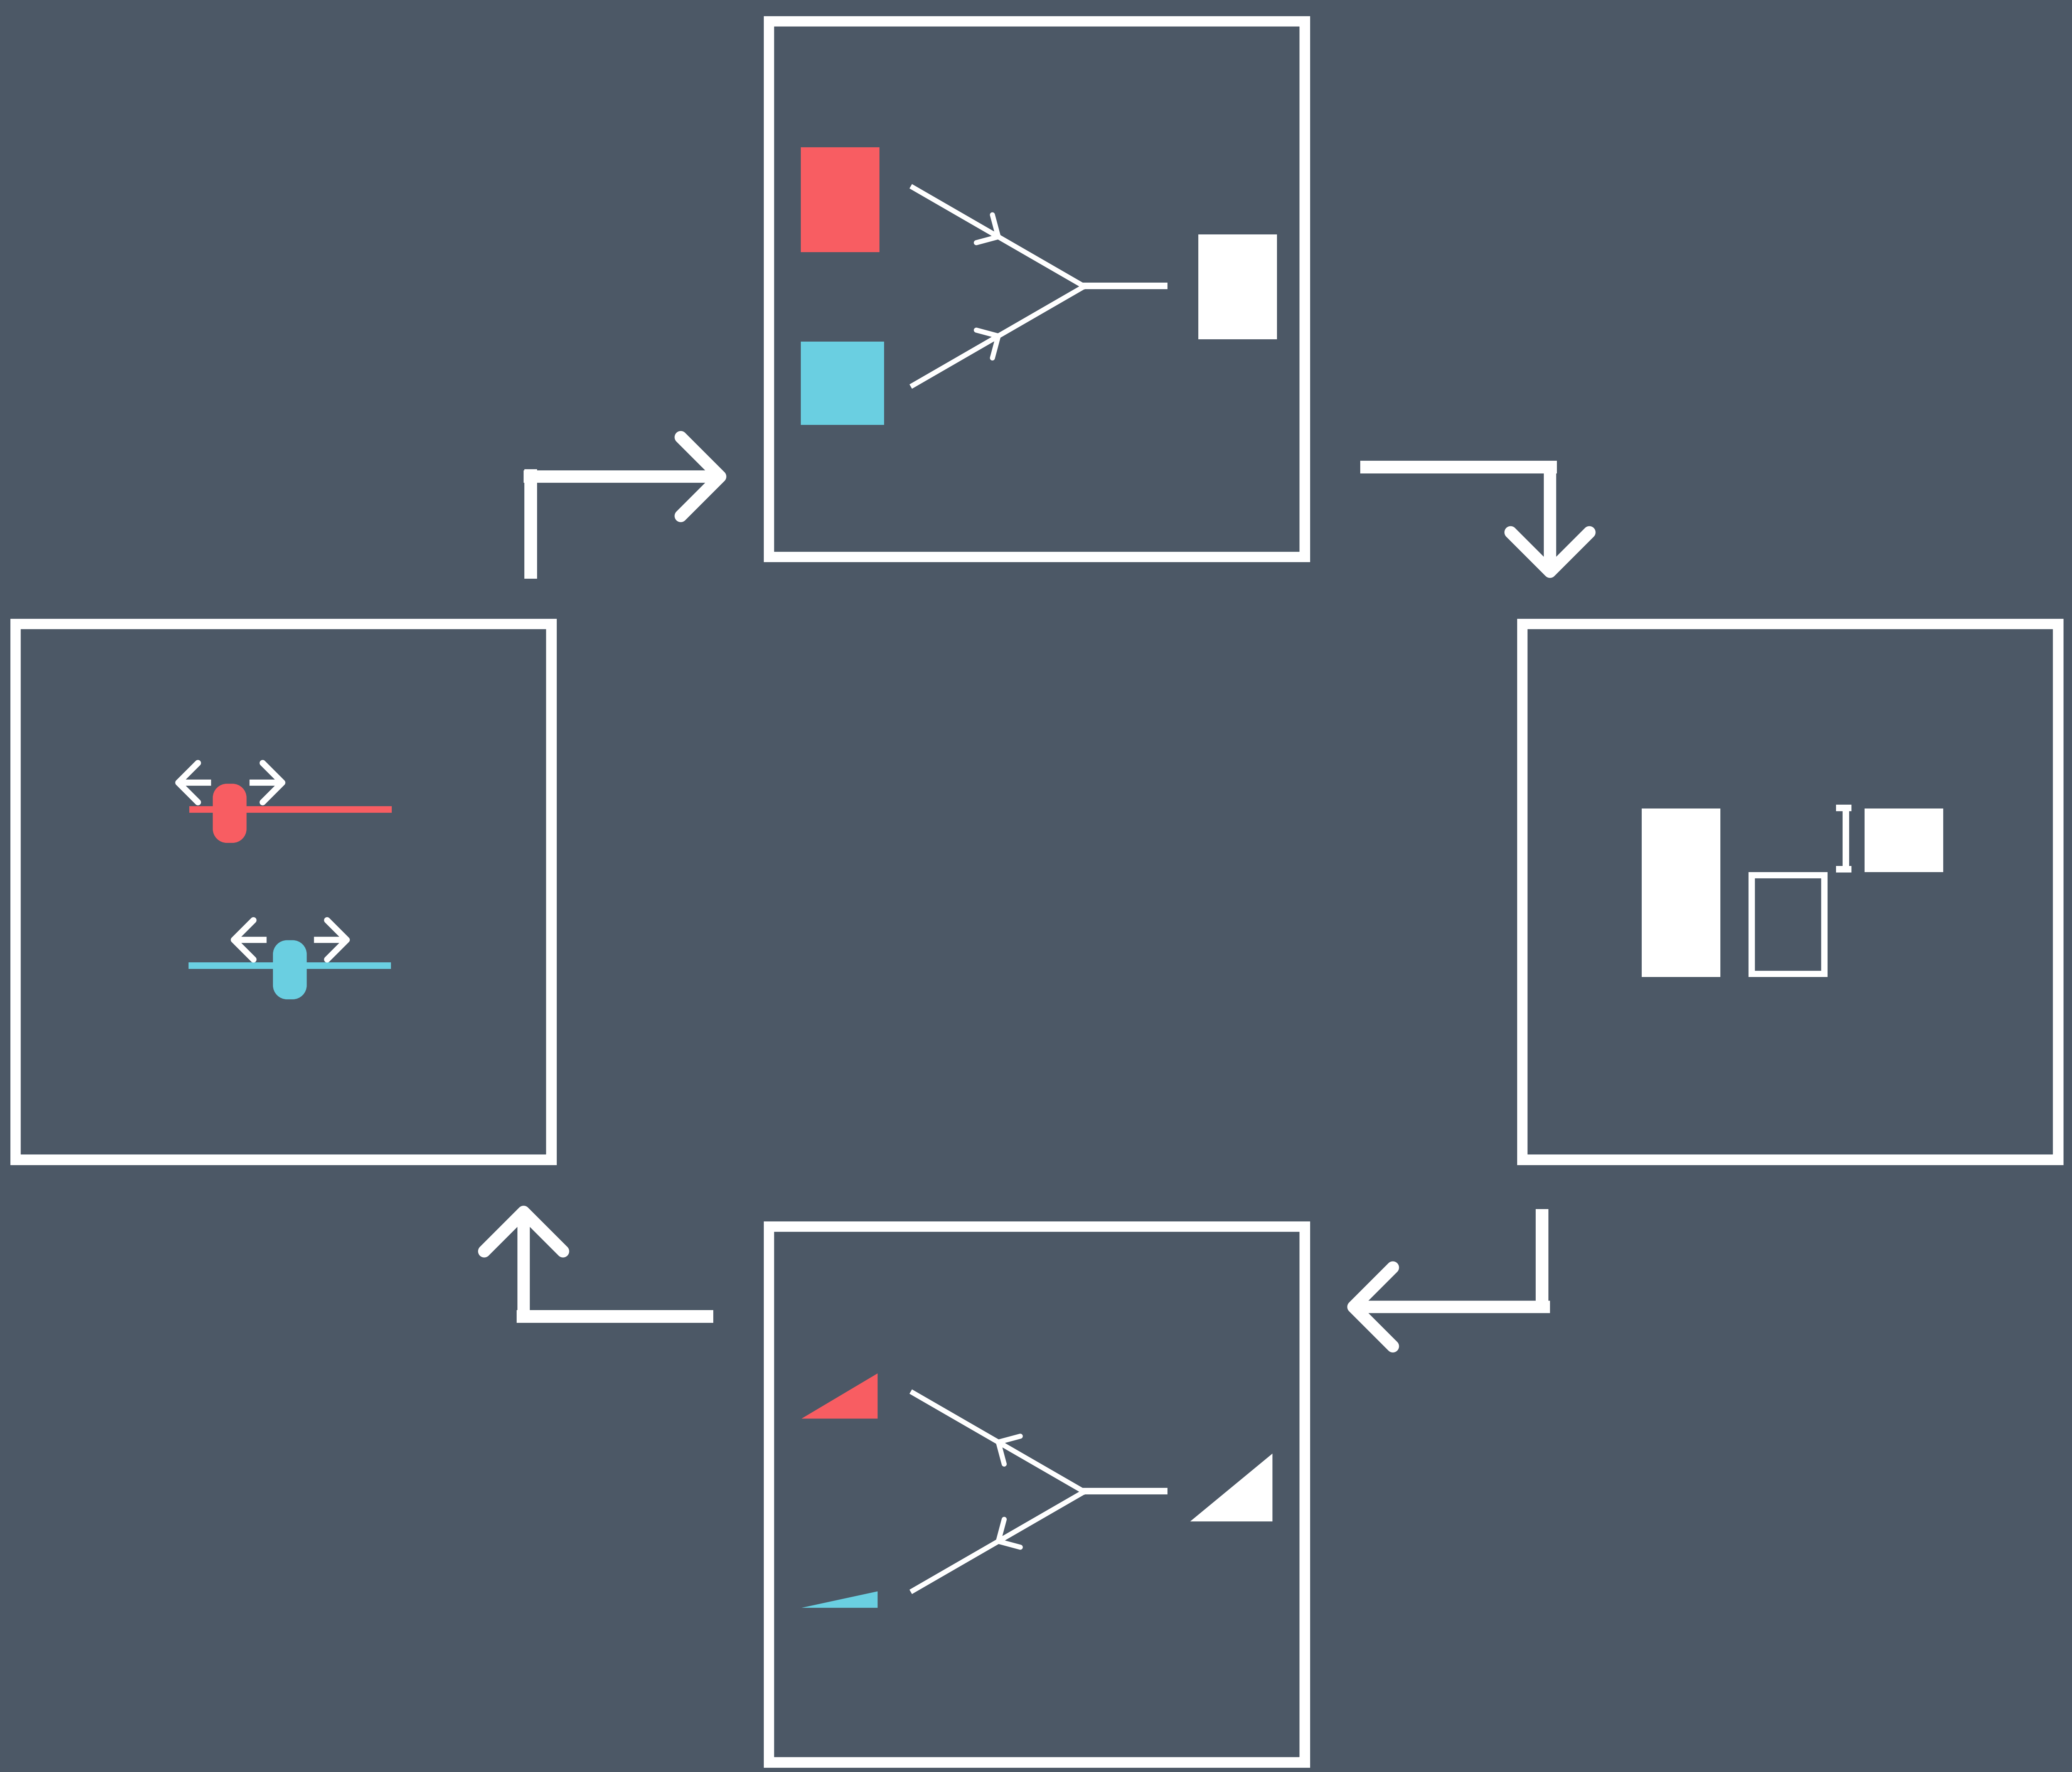

MEOR AMER

ABOUT THIS BOOK

Deep learning is the algorithm powering the current renaissance of artificial intelligence (AI). And its progress is not showing signs of slowing down. A McKinsey report estimates that by 2030, AI will potentially deliver $13 trillion to the global economy, or 16% of the world's current GDP. This opens up exciting career opportunities in the coming decade.

But deep learning can be quite daunting to learn. With the abundance of learning resources in recent years has emerged another problem—information overload.

This book aims to compress this knowledge and make the subject approachable. By the end of this book, you will be able to build a visual intuition about deep learning and neural networks.

WHO SHOULD READ THIS BOOK

If you are new to deep learning, or machine learning in general.

If you already know some background about deep learning but want to gain further intuition.

BOOK FORMAT

This book uses a visuals-first approach. Each page of this book begins with a visual and is supported by concise text.

This book doesn’t include math derivations and code examples. There are some parts where basic math is involved, but generally it is kept to a minimum.

ABOUT THE AUTHOR

My journey into AI began in 2010 after my son was born with a limb difference. I became interested in machine learning for prosthetics and did an MSc at Imperial College London majoring in neurotechnology.

I have also worked in the telecoms data analytics space, serving clients in over 15 countries.

Above all, I am passionate about education and how we learn. I am currently working on projects that explore ways to create alternative learning experiences using visuals, storytelling, and games.

email: contact@kdimensions.com

TABLE OF CONTENTS

INTRODUCTION
- PREDICTION & DECISION 4
- MACHINE LEARNING 6
- DEEP LEARNING 12
- ALGORITHM 13
- DATA 14
- COMPUTATION 15
- ROADMAP 19
- KEY CONCEPTS 20

FOUNDATIONS
- A NEURON 22
- WEIGHTED SUM 24
- WEIGHTS AND BIASES 25
- ACTIVATION 27
- DATA 30
- DATASET 34
- TRAINING 38
- TESTING 40

1 - LINEAR REGRESSION
- INTRODUCTION 43
- GOAL AND DATASET 45
- PREDICT-MEASURE-FEEDBACK-ADJUST 50
- WEIGHTED SUM AND ACTIVATION 52
- LOSS FUNCTION 58
- MEAN SQUARED ERROR 59
- MINIMIZING LOSS 64
- GRADIENT 68
- GRADIENT DESCENT 73
- LEARNING RATE 76
- EPOCH 81
- COST AND METRIC 83
- PERFORMANCE 86

2 - NON-LINEAR REGRESSION
- INTRODUCTION 93
- GOAL AND DATASET 99
- ARCHITECTURE 102
- PREDICT 108
- MEASURE 112
- FEEDBACK 114
- COMPUTATION GRAPH 117
- BACKPROPAGATION 119
- ADJUST 129
- PERFORMANCE 135
- LINEAR ACTIVATION 137
- LINEARITY 139
- NON-LINEARITY 141
- RELU ACTIVATION 142
- PERFORMANCE 145
- ACTIVATION FUNCTIONS 148

3 - BINARY CLASSIFICATION
- INTRODUCTION 150
- CLASSIFICATION VS. REGRESSION 152
- GOAL AND DATASET 153
- ARCHITECTURE 156
- SIGMOID ACTIVATION 158
- BINARY CROSS ENTROPY 166
- ACCURACY 169
- PERFORMANCE 170
- CONFUSION MATRIX 174
- PRECISION-RECALL 176
- F1 SCORE 177

4 - MULTI-CLASS CLASSIFICATION
- INTRODUCTION 180
- GOAL AND DATASET 181
- ONE-HOT ENCODING 183
- ARCHITECTURE 184
- SOFTMAX ACTIVATION 188
- CATEGORICAL CROSS ENTROPY 195
- PERFORMANCE 197
- IMPROVING PERFORMANCE
  - HYPERPARAMETERS 202
  - DATA TECHNIQUES 210

THE BIGGER PICTURE
- NEURAL NETWORKS
  - FEEDFORWARD 217
  - CONVOLUTIONAL 219
  - RECURRENT 225
  - GENERATIVE ADVERSARIAL 230
  - OTHER ARCHITECTURES 232
- CONCLUSION 233
- KEY CONCEPTS REVISITED 234
- SUGGESTED RESOURCES 235

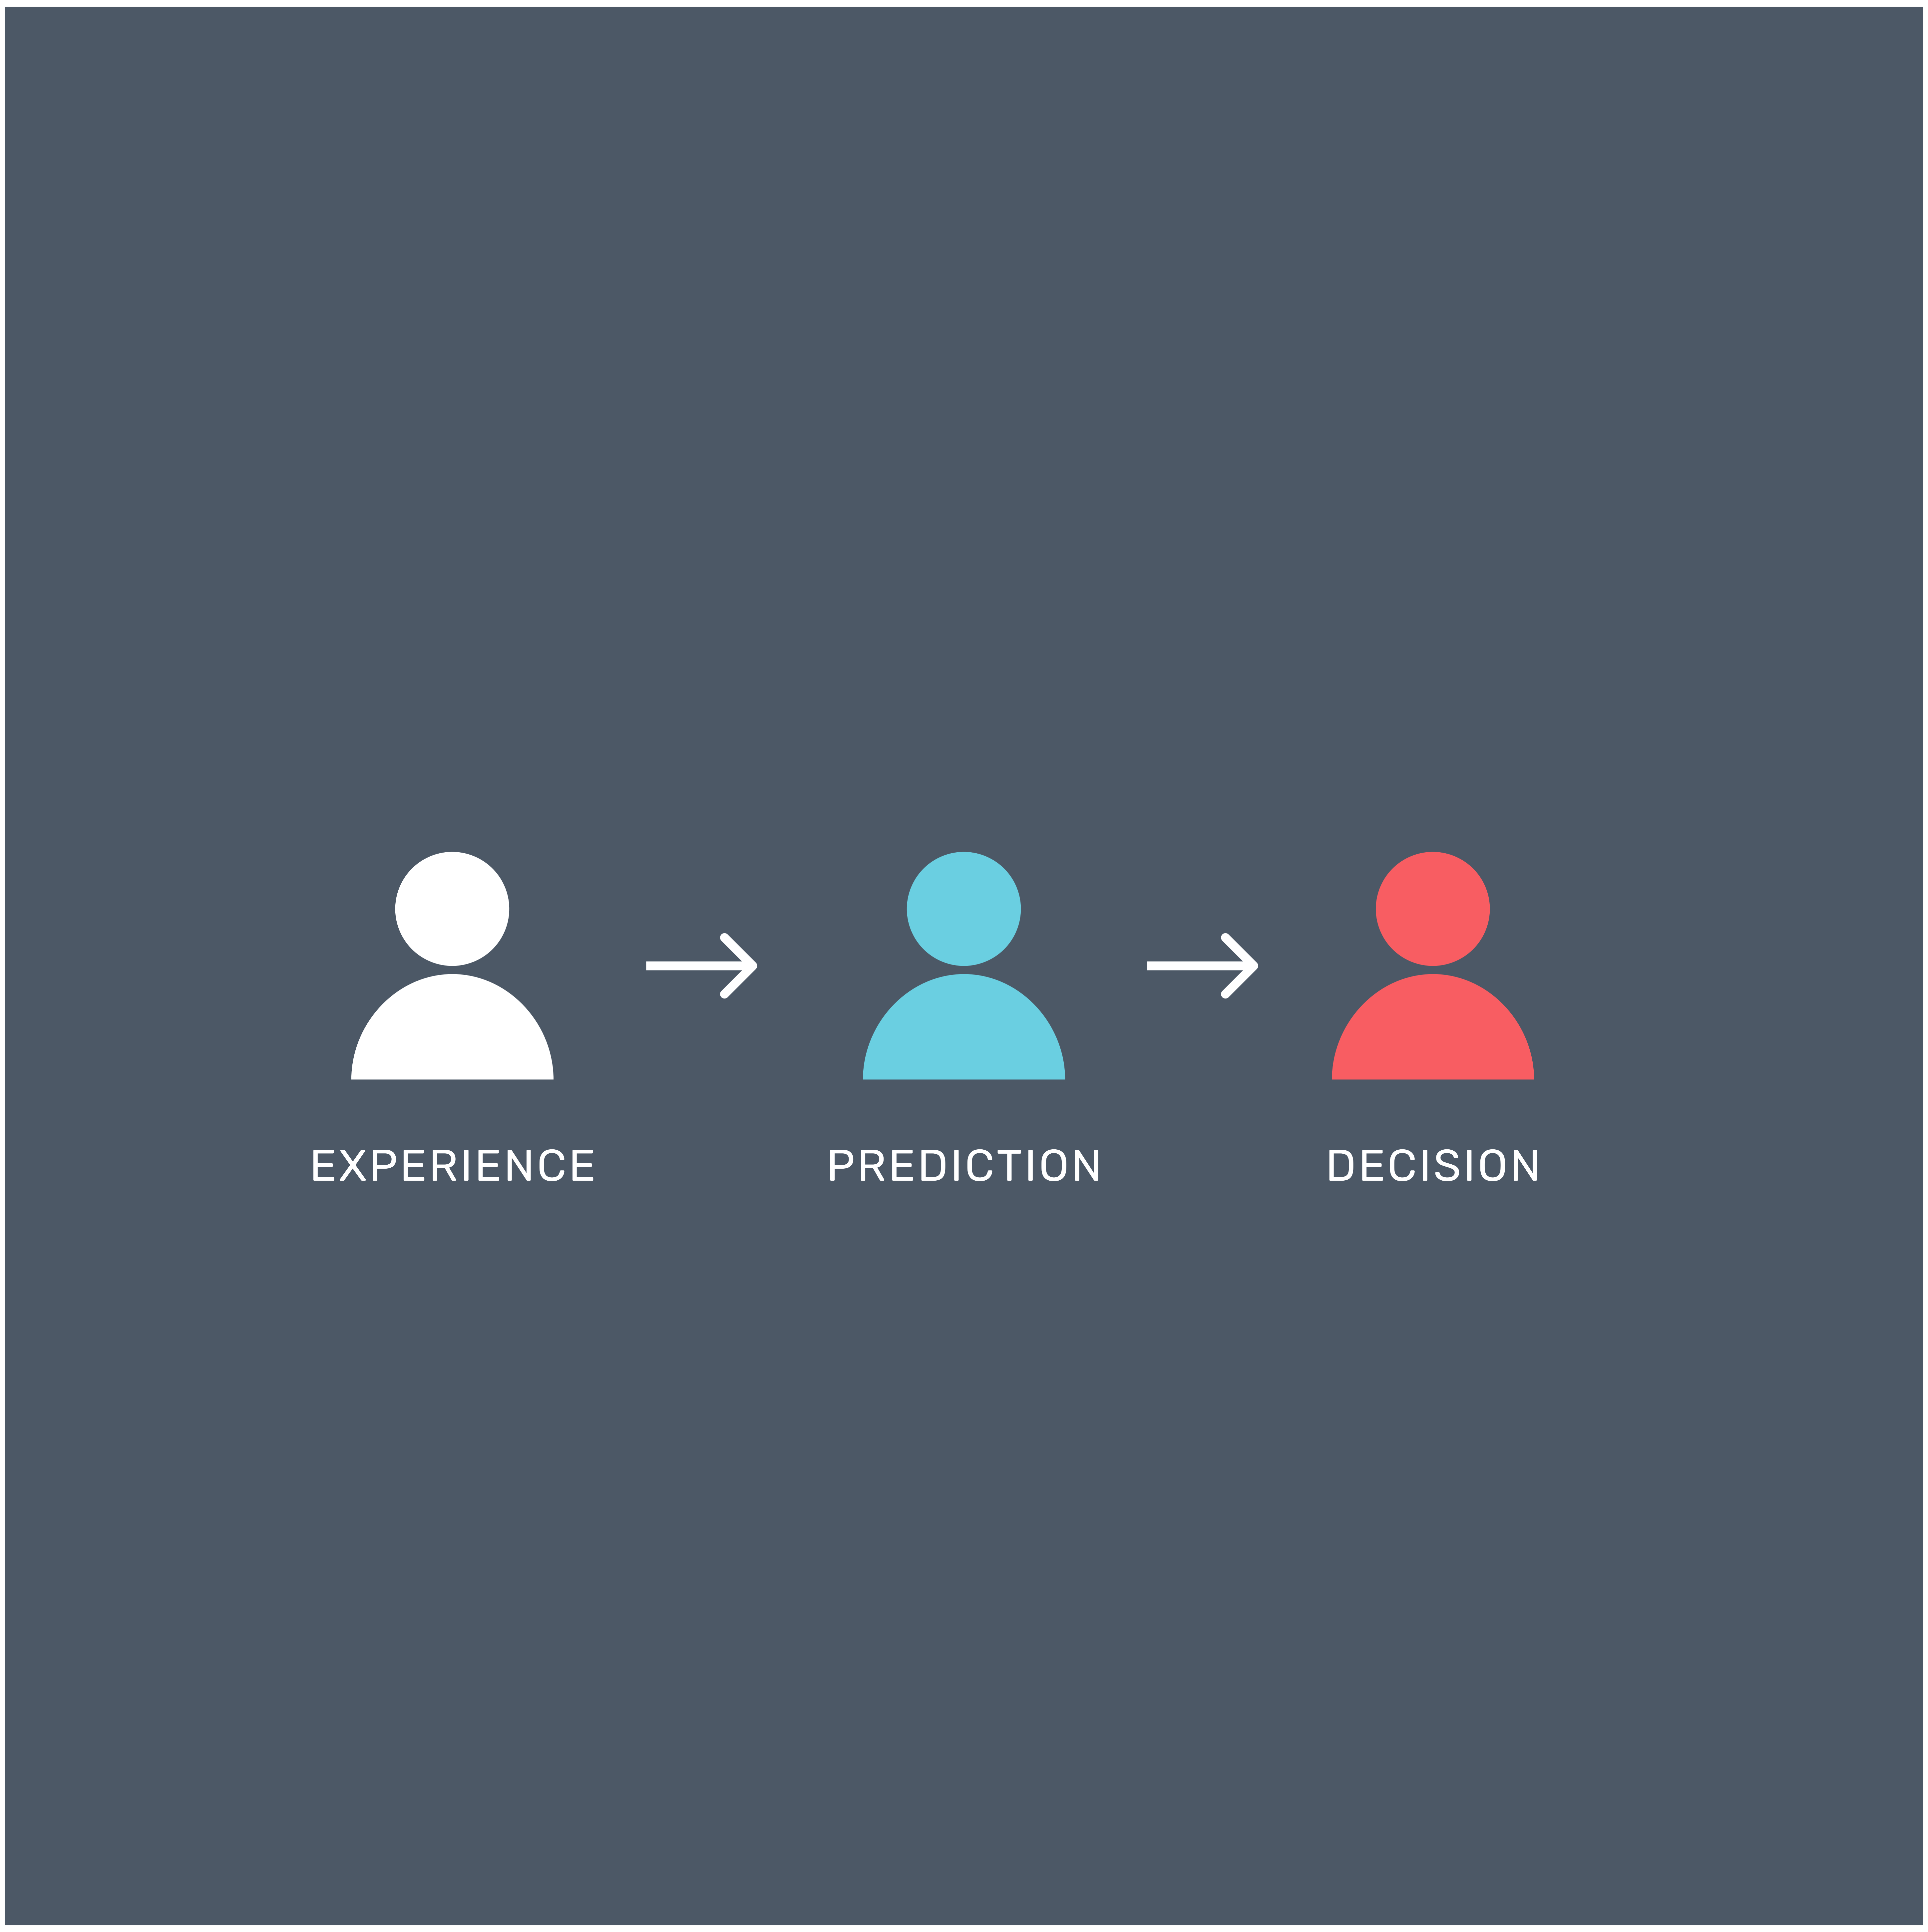

PREDICTION AND DECISION

Prediction is a key ingredient in decision-making under uncertainty. — Prediction Machines book.

Much that goes on in our lives involves some form of prediction. These predictions differ in one way, namely, how sure we are of them. In some tasks, they don't feel like predictions because we feel so sure about them. In some others, we know next to nothing about them, so they become mere guesses.

All of this depends on how simple a task is and, more importantly, how much experience we have with it.

To illustrate this, let's look at some examples.

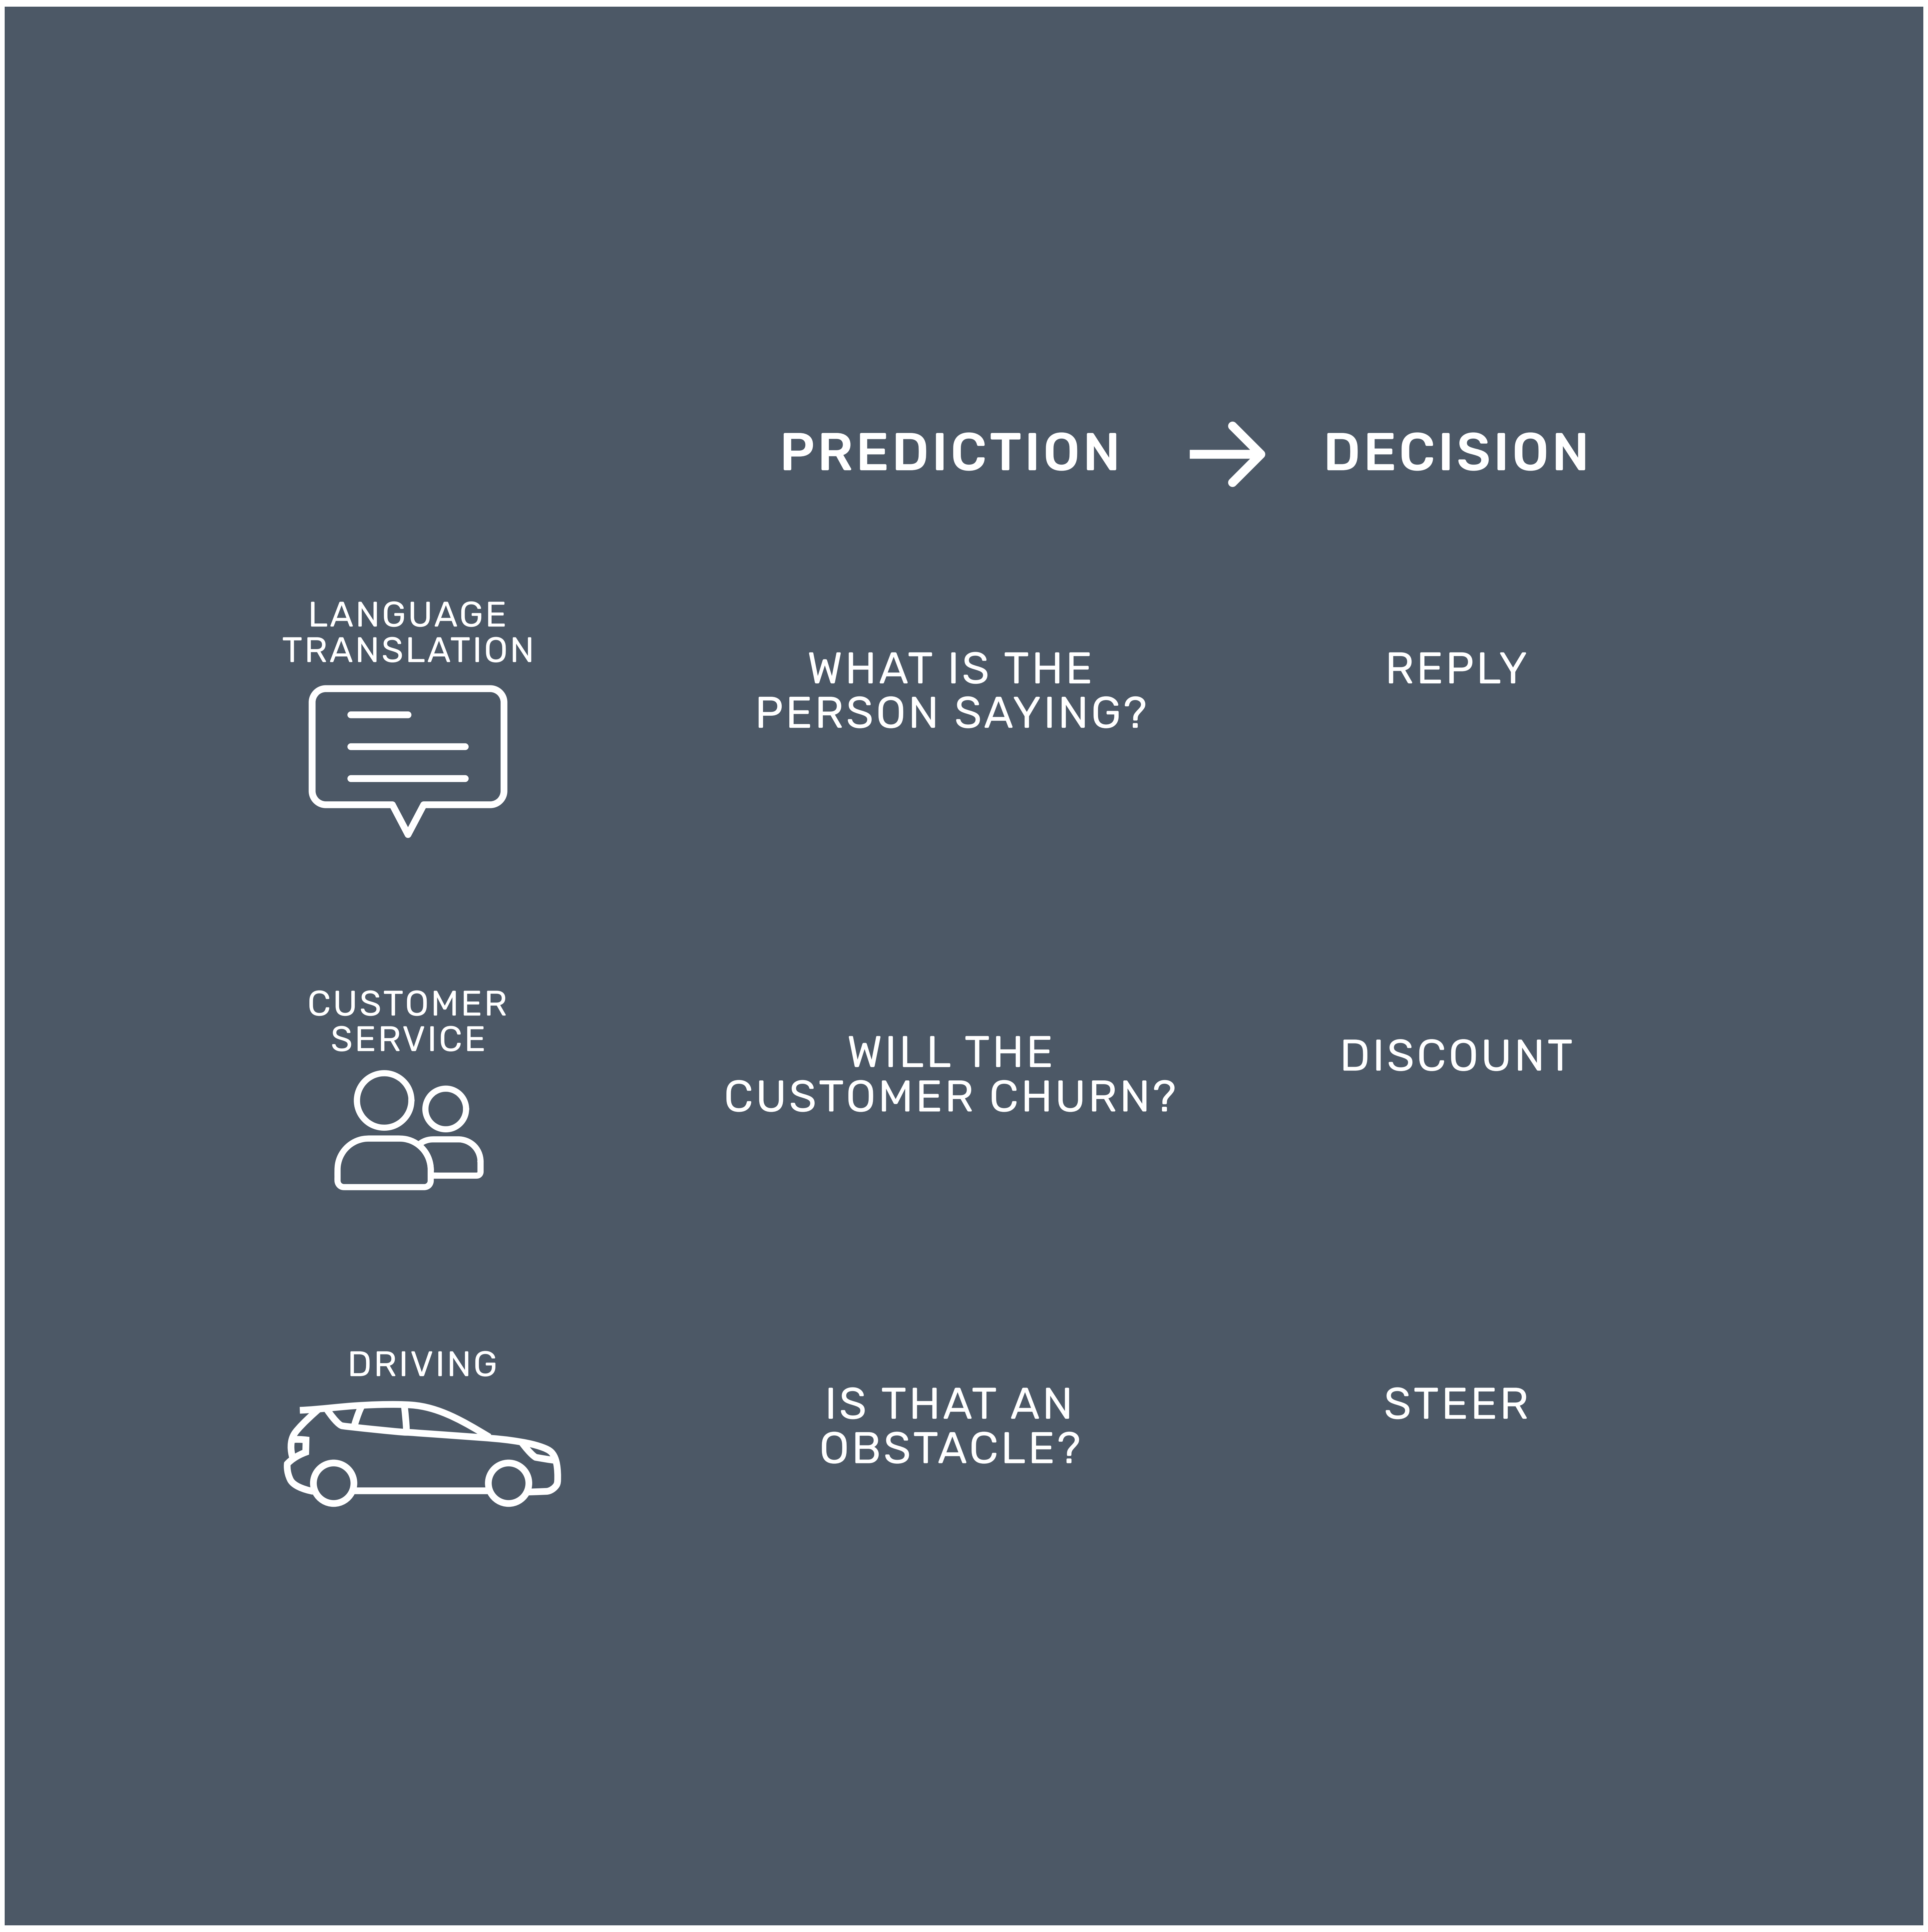

EXAMPLES

Let's take the example of language translation. As we listen to someone speaking, we are predicting what the person means. The more experience we have with this language, the better our prediction becomes, and the better our decision, that is our reply, becomes.

Take another example in a business setting. Our experience dealing with customers can help us see patterns in their behavior, so we'll notice if they are likely to churn.

As for driving, the more miles we clock, the more skilled we become and the more adept we are at evaluating our surroundings.

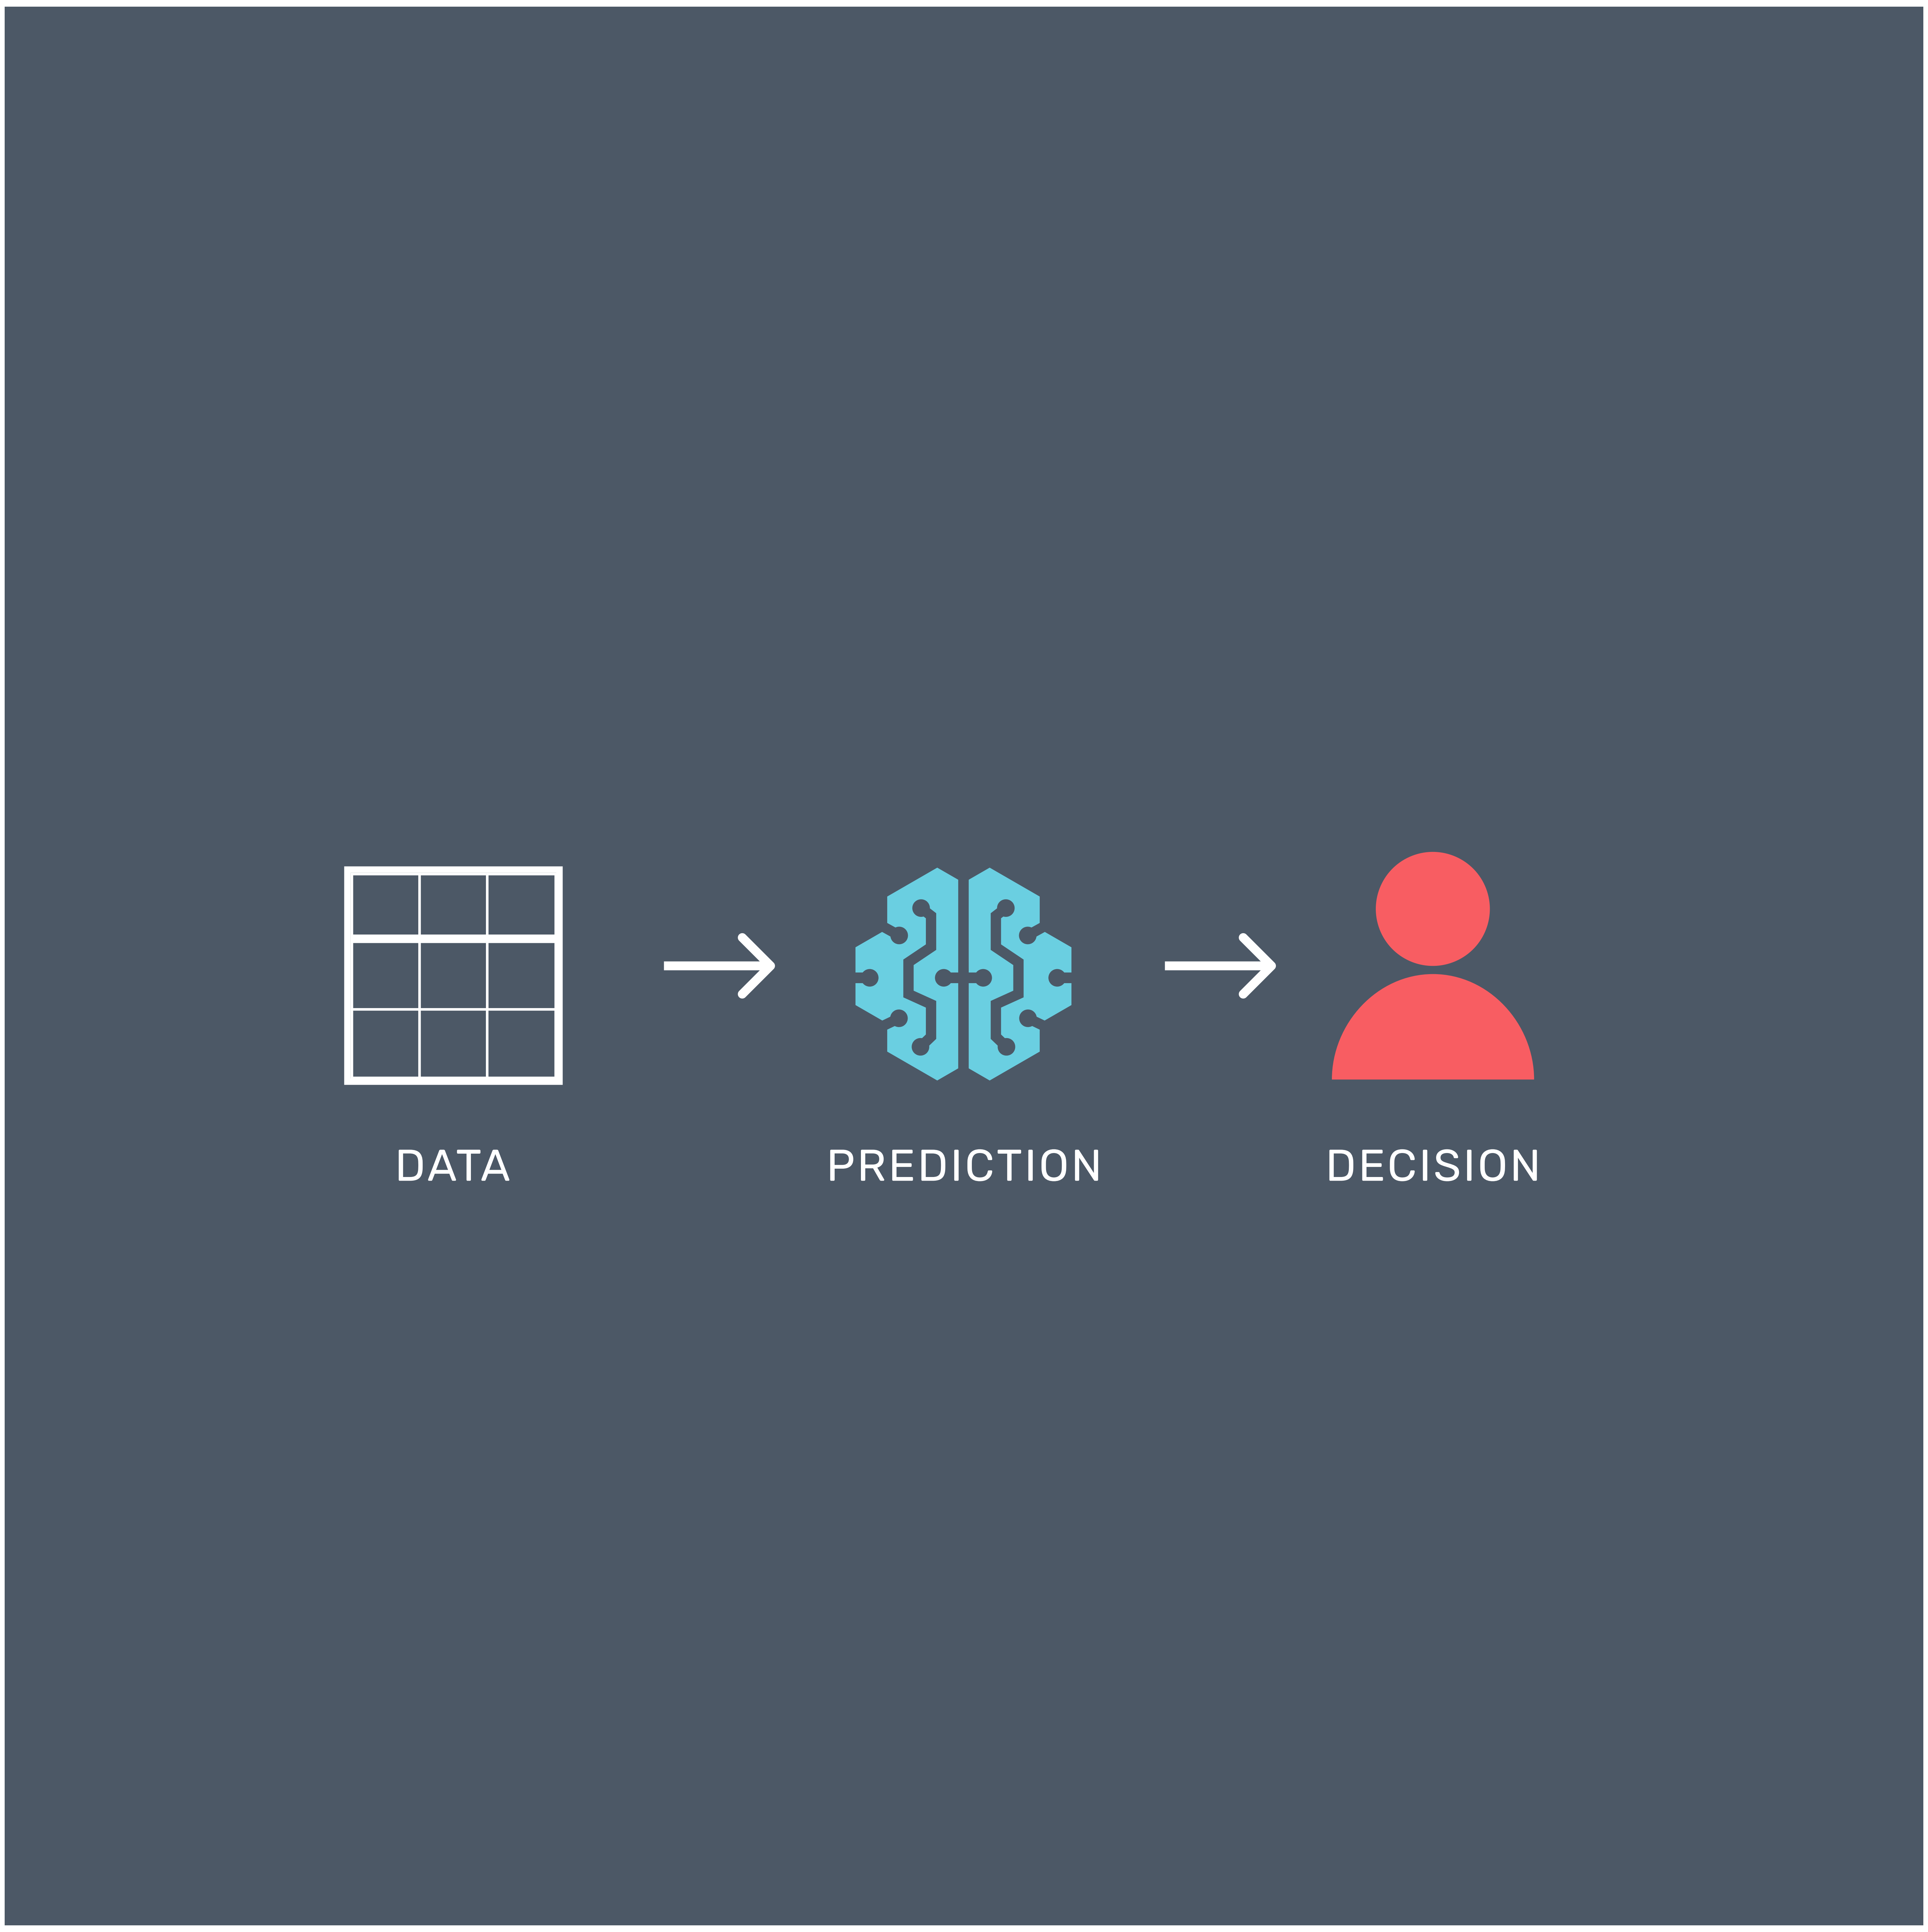

WHAT IS MACHINE LEARNING?

In many of these tasks, machine learning can handle the prediction on our behalf.

In recent years, the adoption of machine learning has accelerated. Many industries and verticals are already deploying use cases that automate predictions using machine learning.

In the machine's world, the experience comes in the form of data. Just as we learn from experience, the machine learns from data.

That is what machine learning is all about—learning from the data and turning it into predictions.

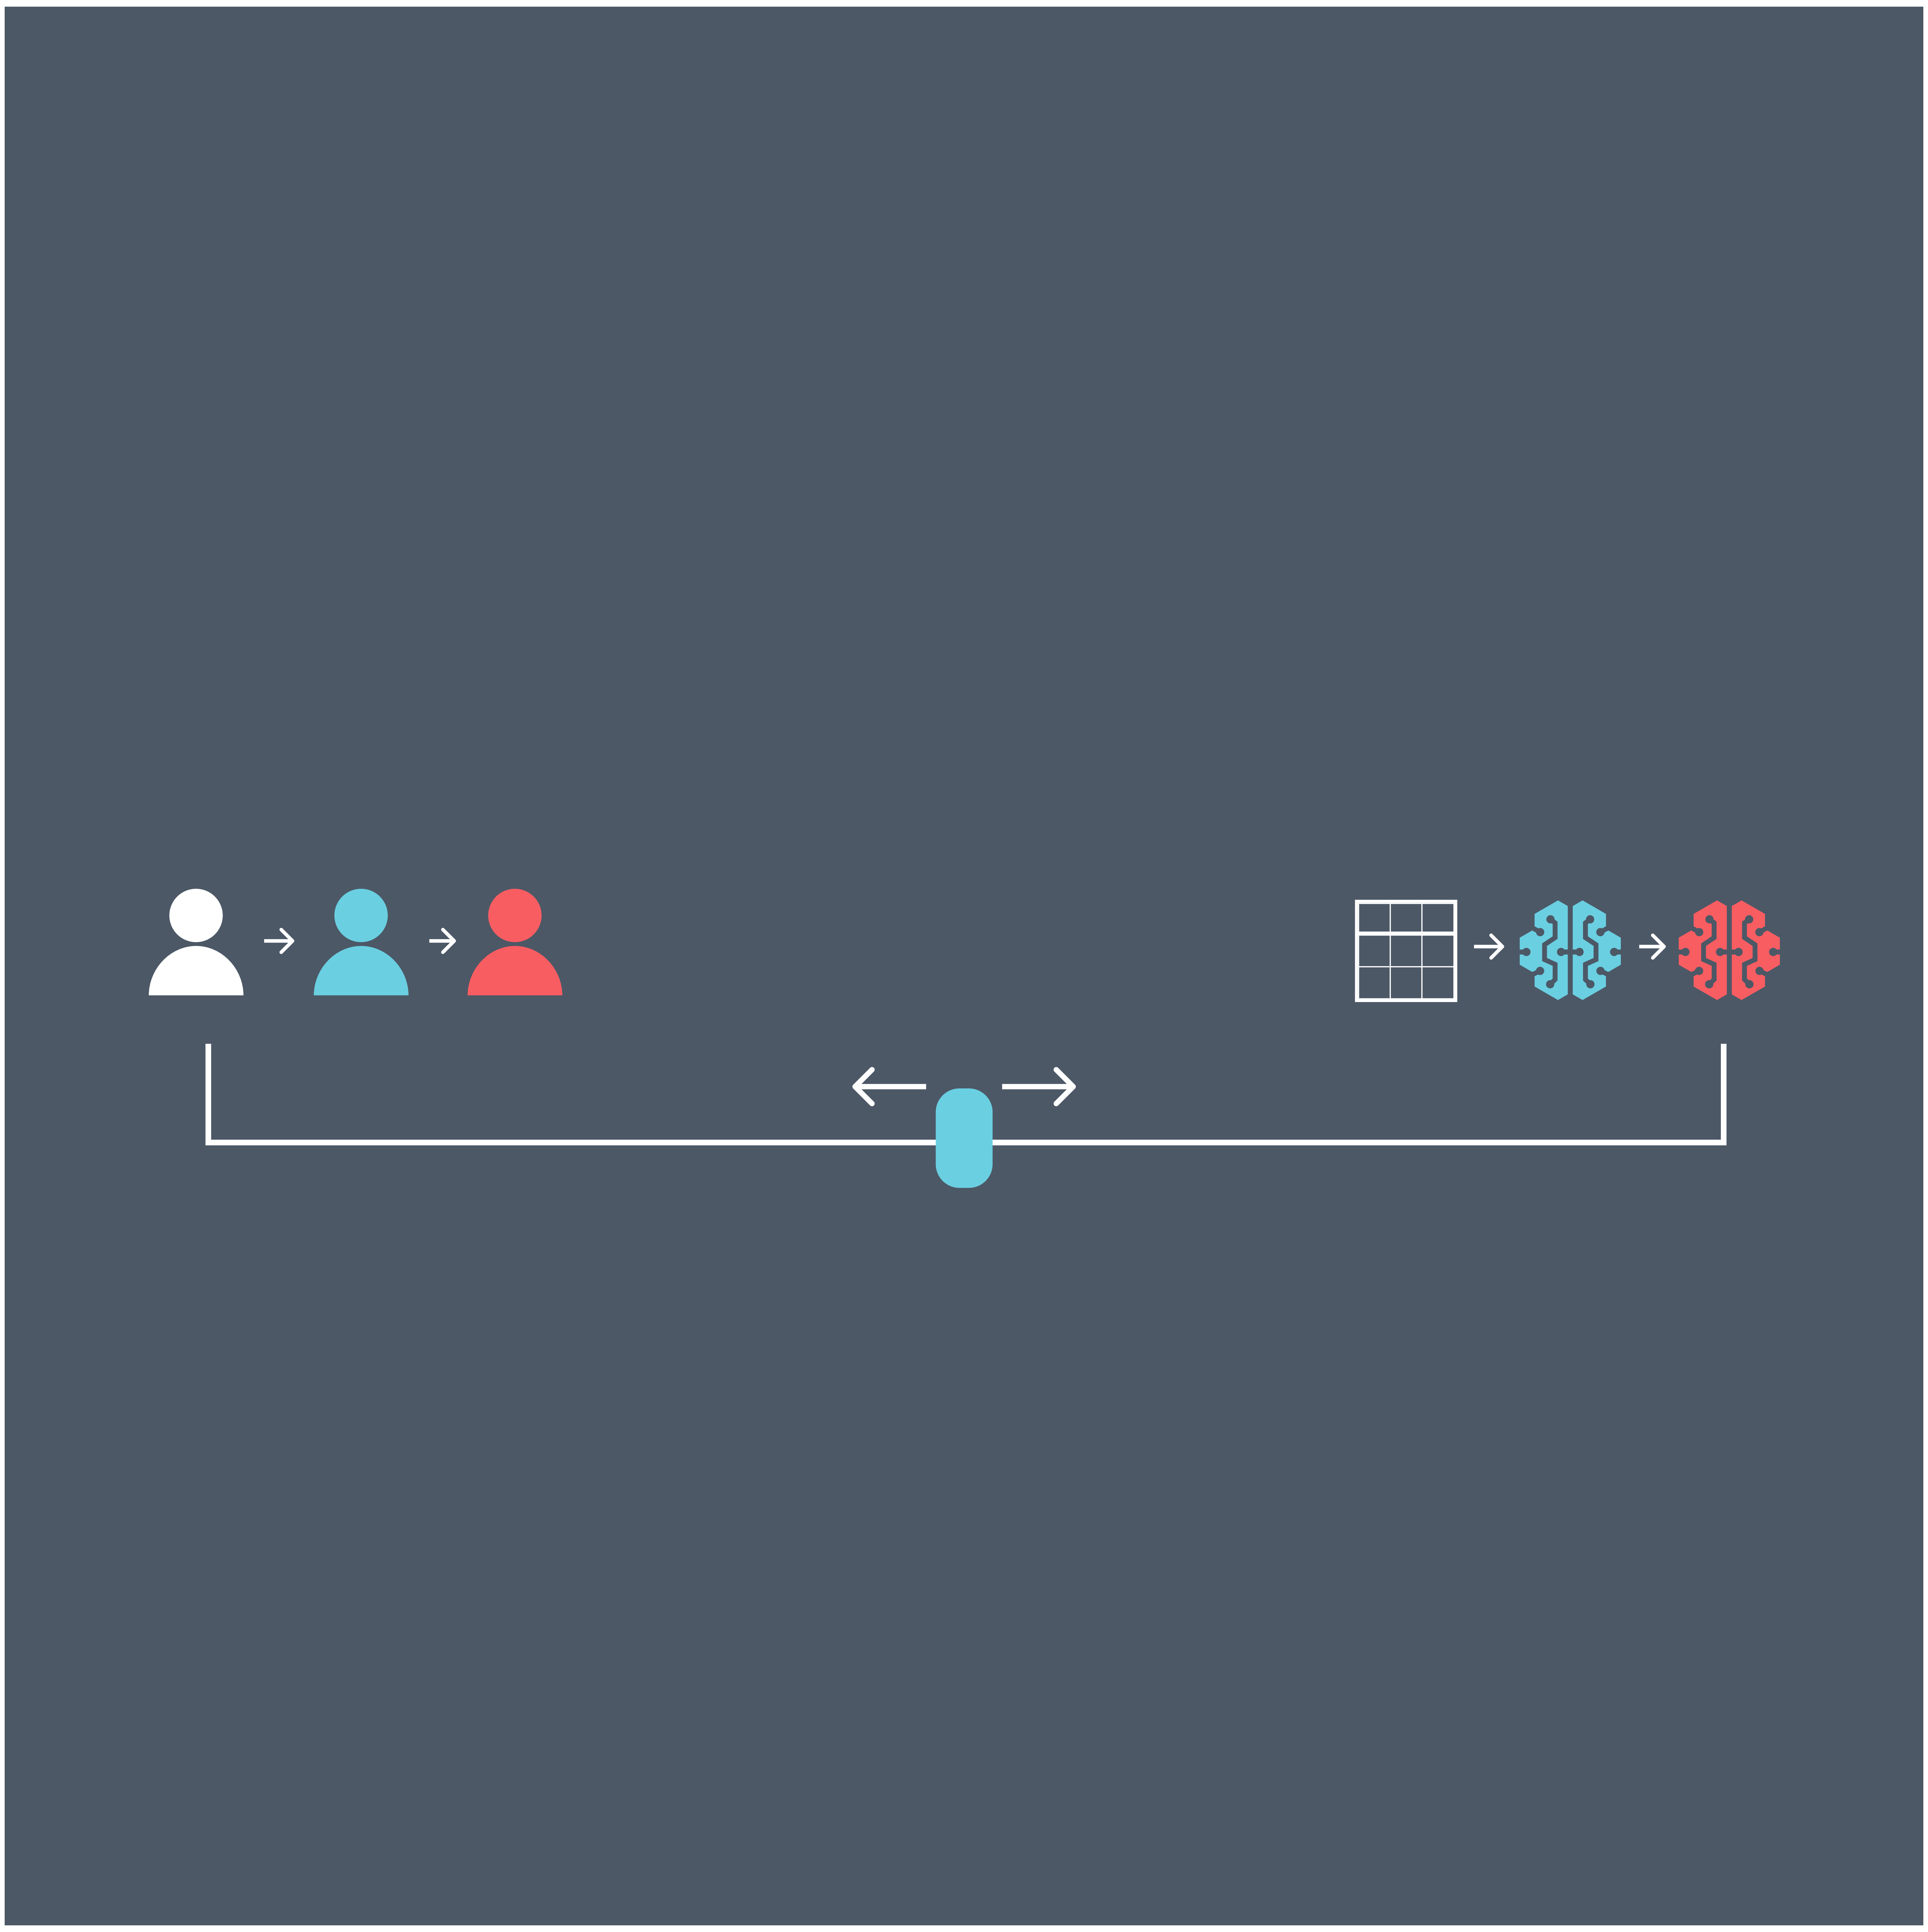

MACHINE LEARNING IN THE REAL WORLD

In fact, machine learning can even handle the decision part. In some domains, most notably self-driving cars, we are not far from seeing full automation becoming the norm.

But in most other domains, this is still far from reality. For this reason, the focus of this book is on the prediction part.

And indeed, it is upon us to ensure healthy technological progress, where people can thrive with the help of machines rather than being inhibited by them. That's the sweet spot that we are collectively trying to find.

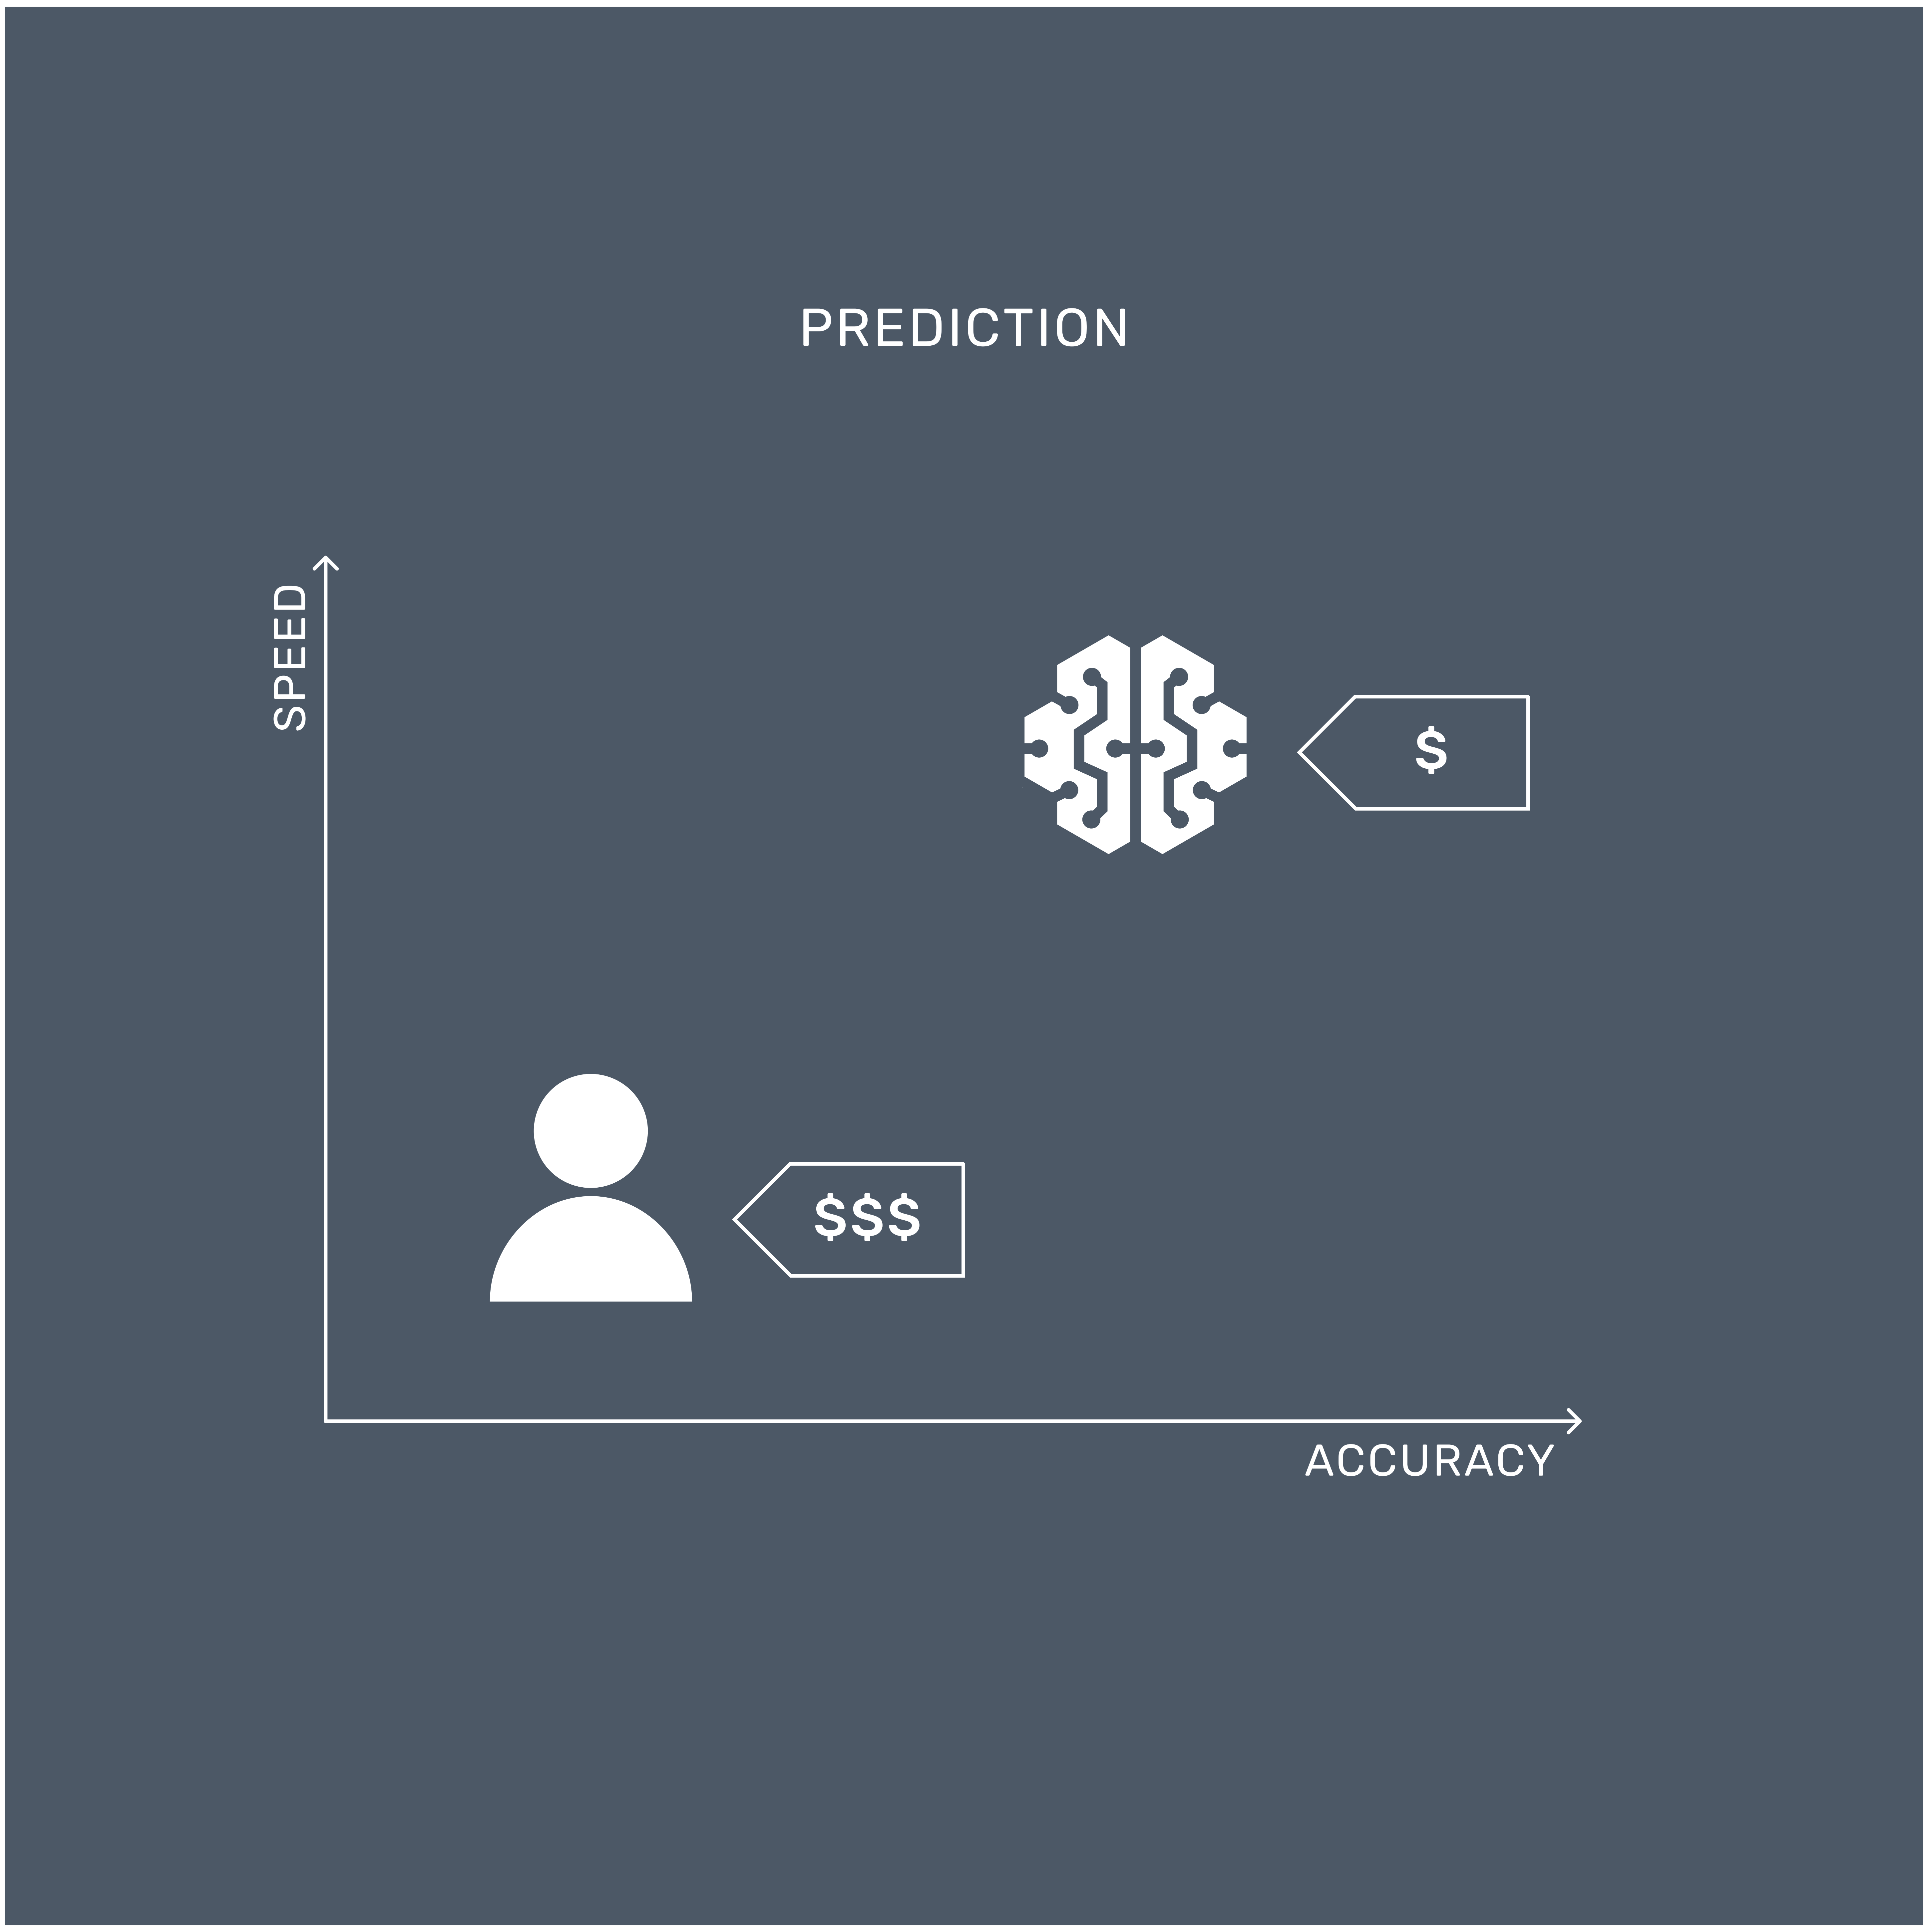

THE VALUE OF MACHINE LEARNING

So, what is the value of having machines that can make predictions on our behalf?

In the book Prediction Machines, the authors argued for a few reasons why prediction machines are so valuable, the first being that ‘they can often produce better, faster, and cheaper predictions than humans can’.

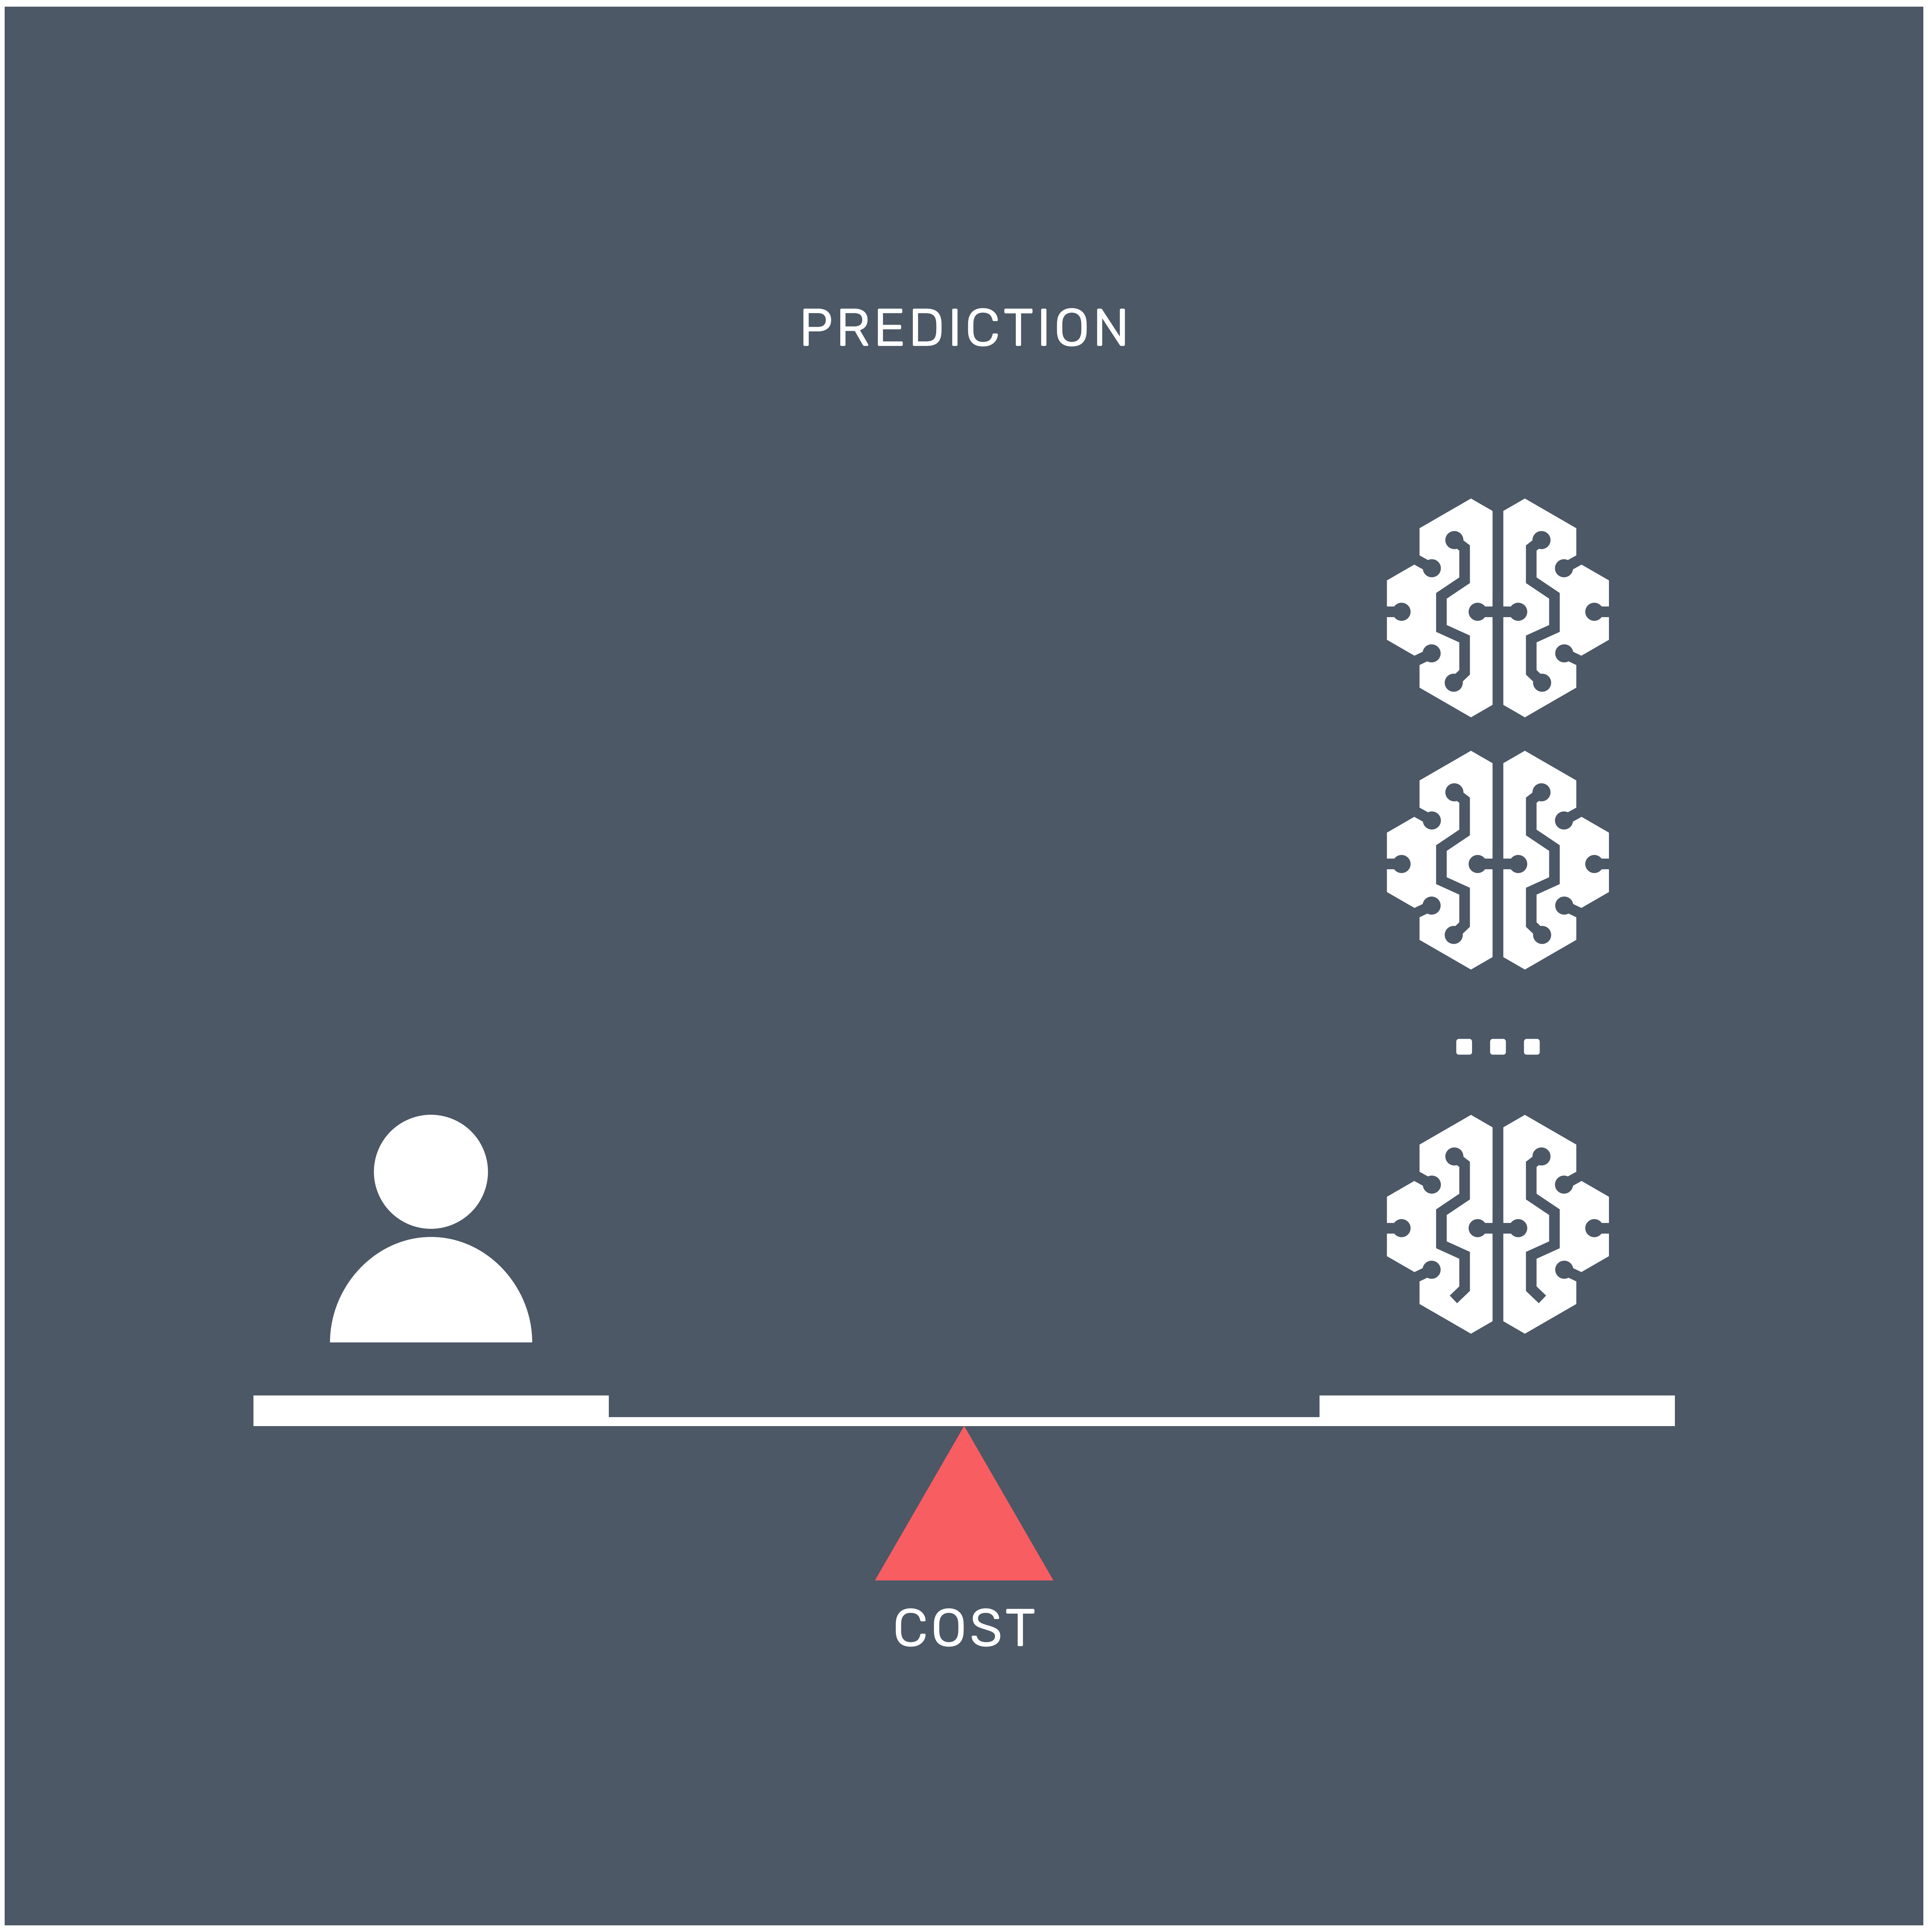

ACCELERATING HUMAN PROGRESS

The cheaper the cost of prediction, the more tasks we can take on. The world is full of challenges waiting to be solved. Machine learning enables us to scale our efforts in ways that have not been possible before, presenting us with the opportunity to take on these challenges.

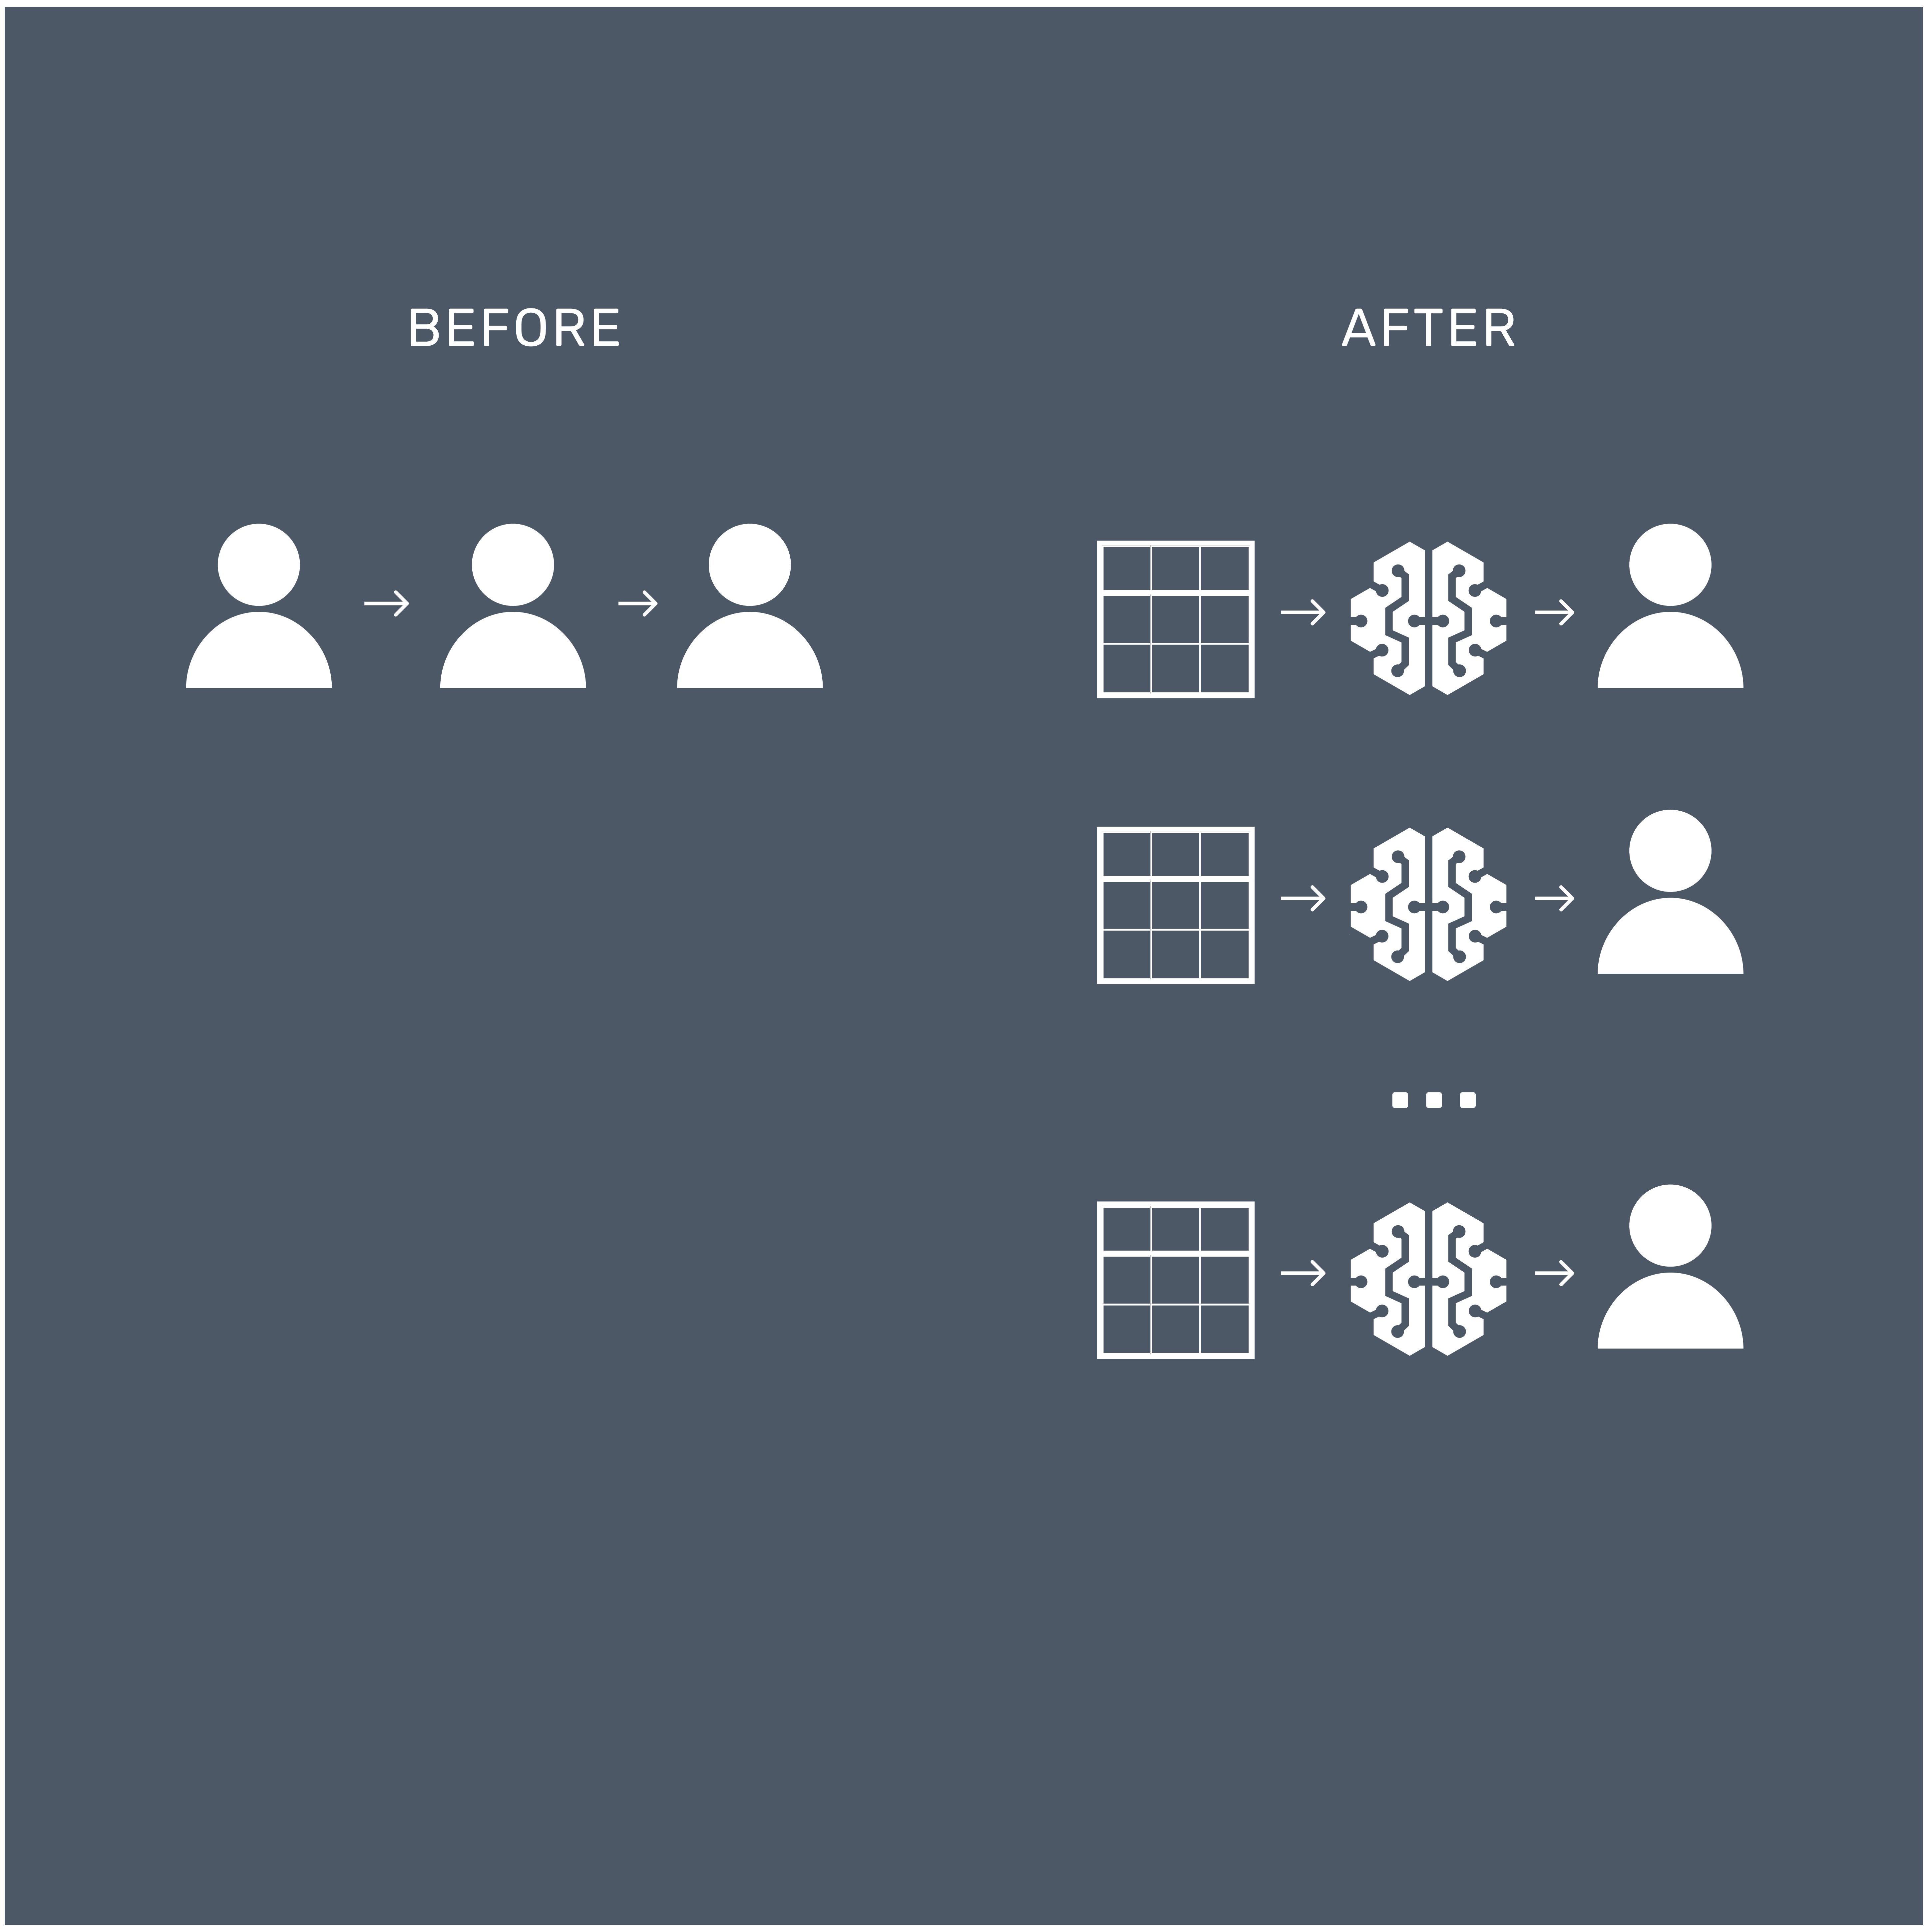

EVOLUTION IN ROLES

Some may worry that this will spell the end of most jobs, and rightly so. But looking at the bigger picture, there will in fact be even more job opportunities.

The World Economic Forum's The Future of Jobs Report 2020 estimates that by 2025, 85 million jobs may be displaced. But on the other hand, 97 million new roles may emerge. This already takes into account the economic slowdown due to the pandemic, and still, the net effect is positive.

Job roles will evolve, and the machine's role is to serve us so we can pursue more creative and challenging endeavors.

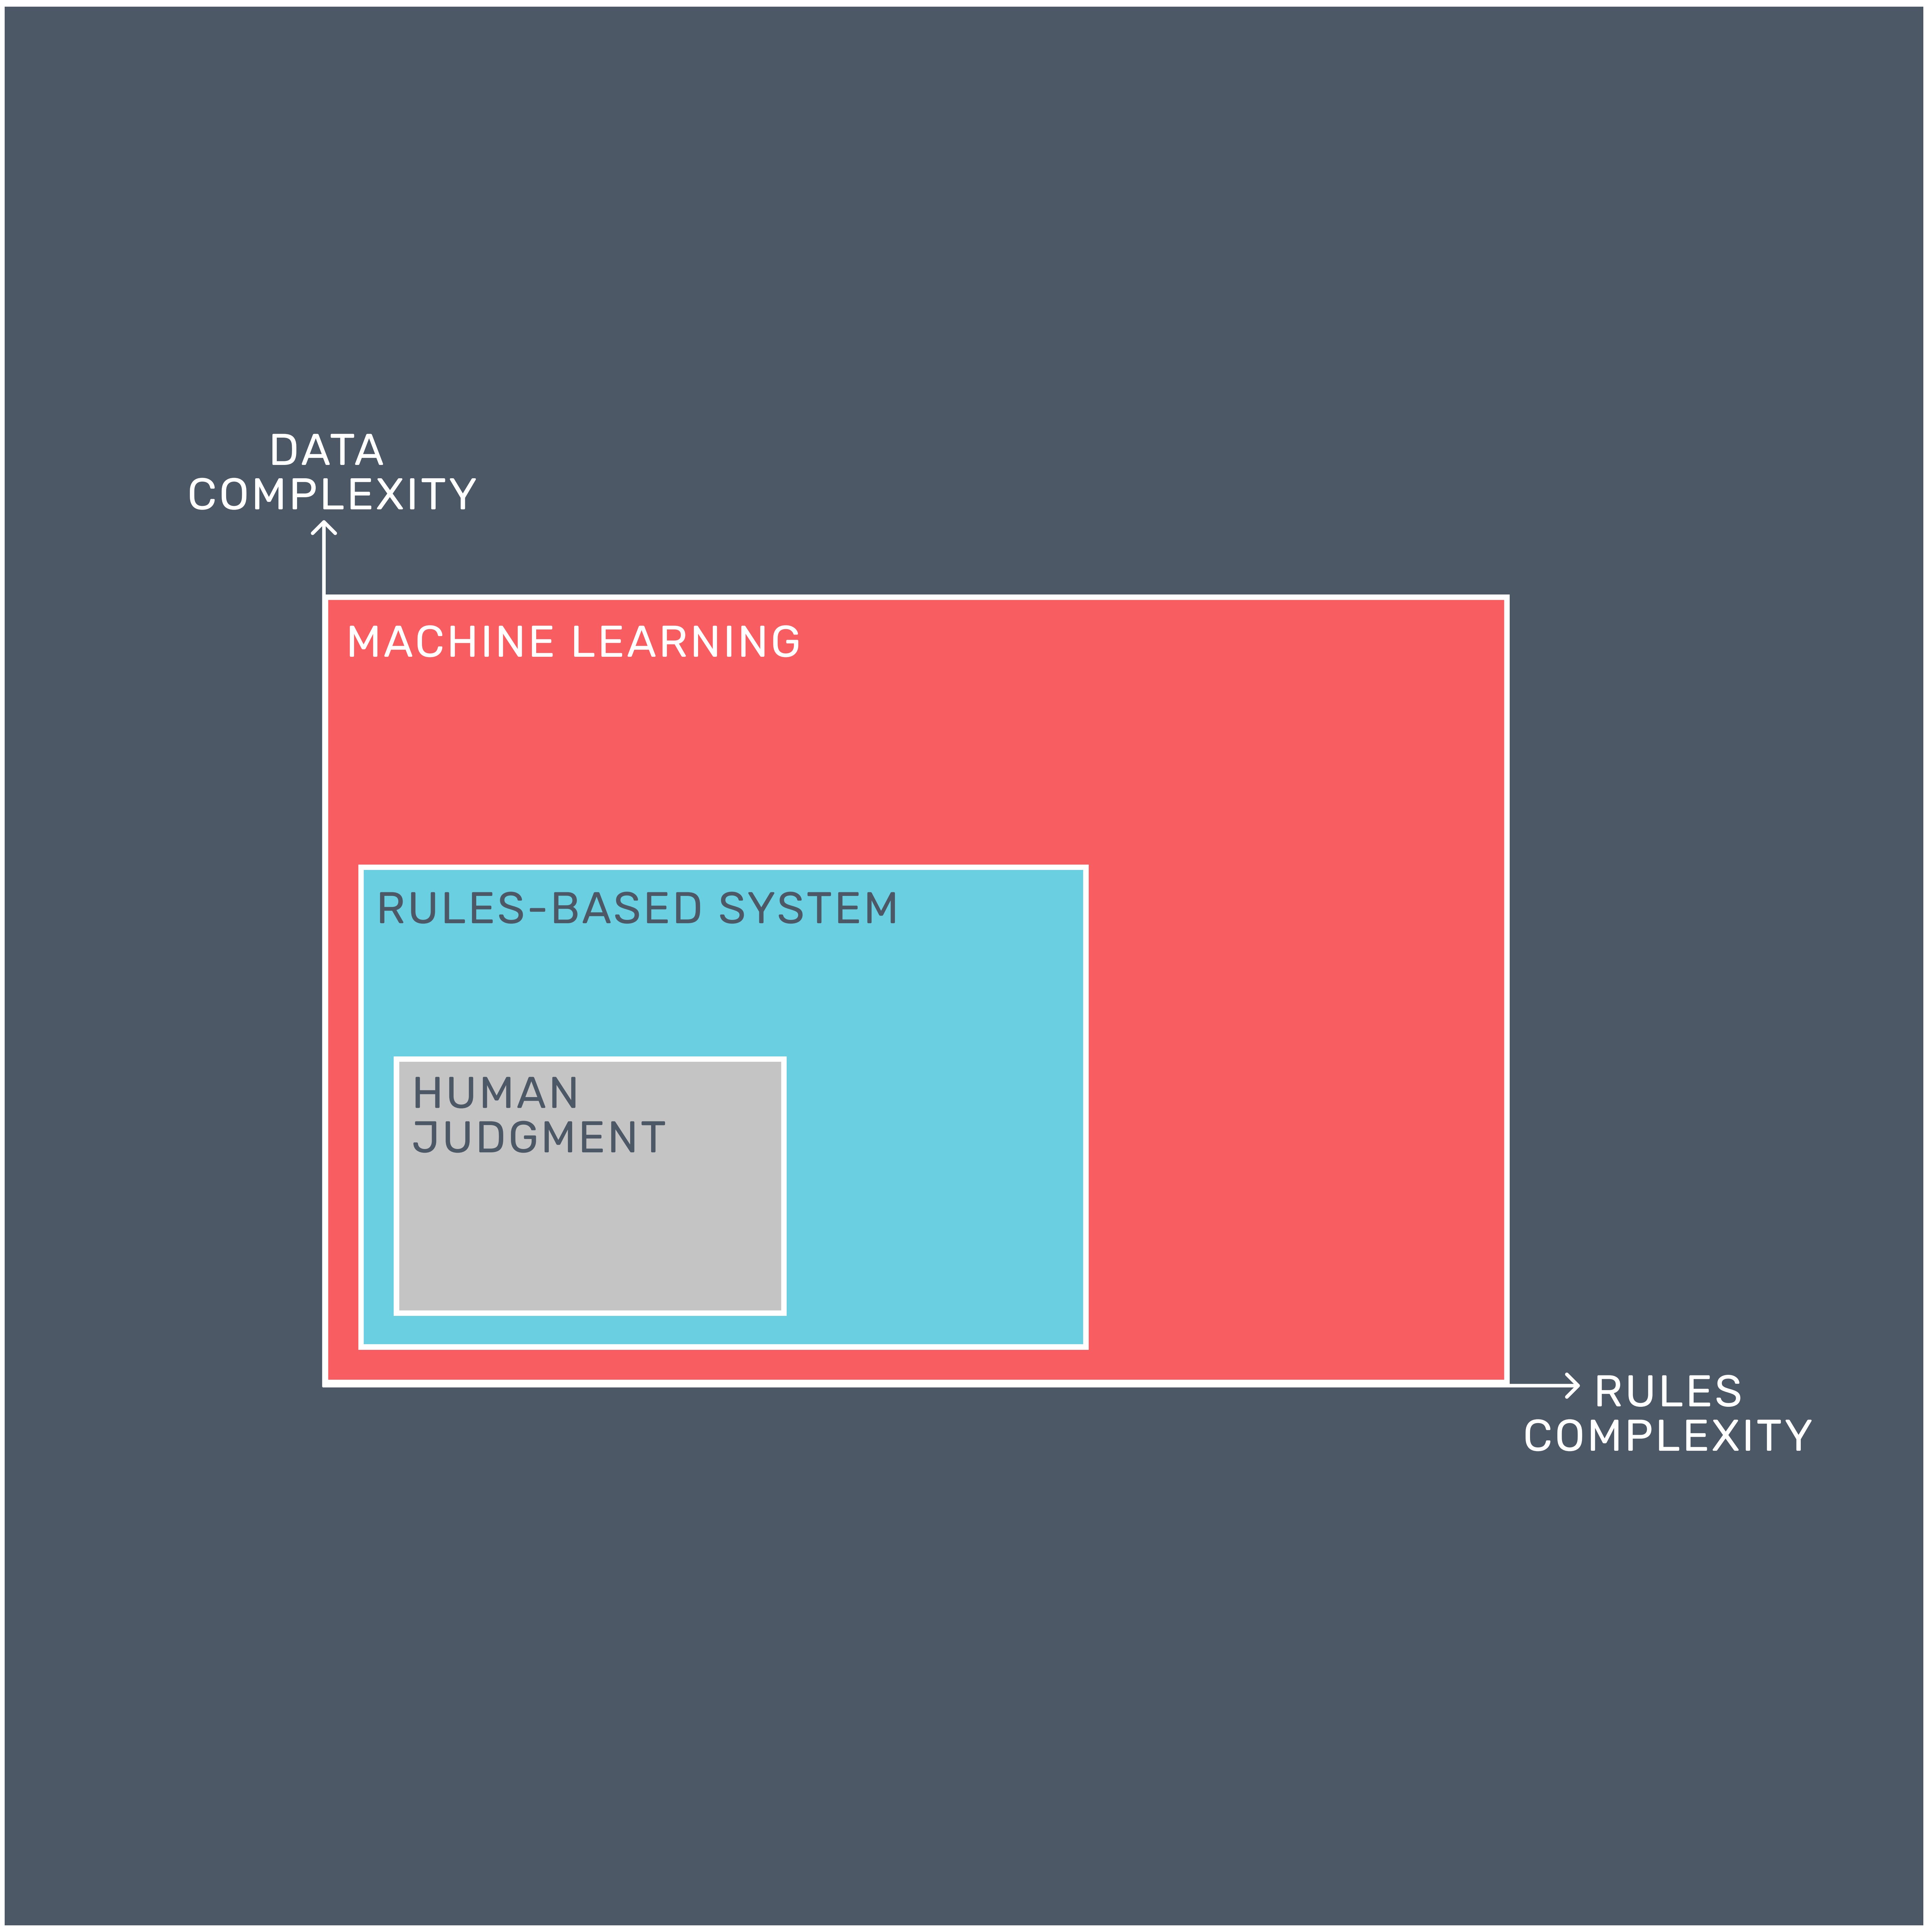

WHEN TO USE MACHINE LEARNING?

We can think of prediction automation in three phases.

The first, that is without automation, is relying on human judgment, either based on data or experience.

The second is using a rules-based system. We translate our experience into rules that software can understand and execute based on data as inputs.

The third is machine learning, which uses data to create its own rules, guided by the goal defined by humans.

As the data and rules become more complex, it makes sense to use machine learning. Otherwise, it may not be cost-effective to do so.

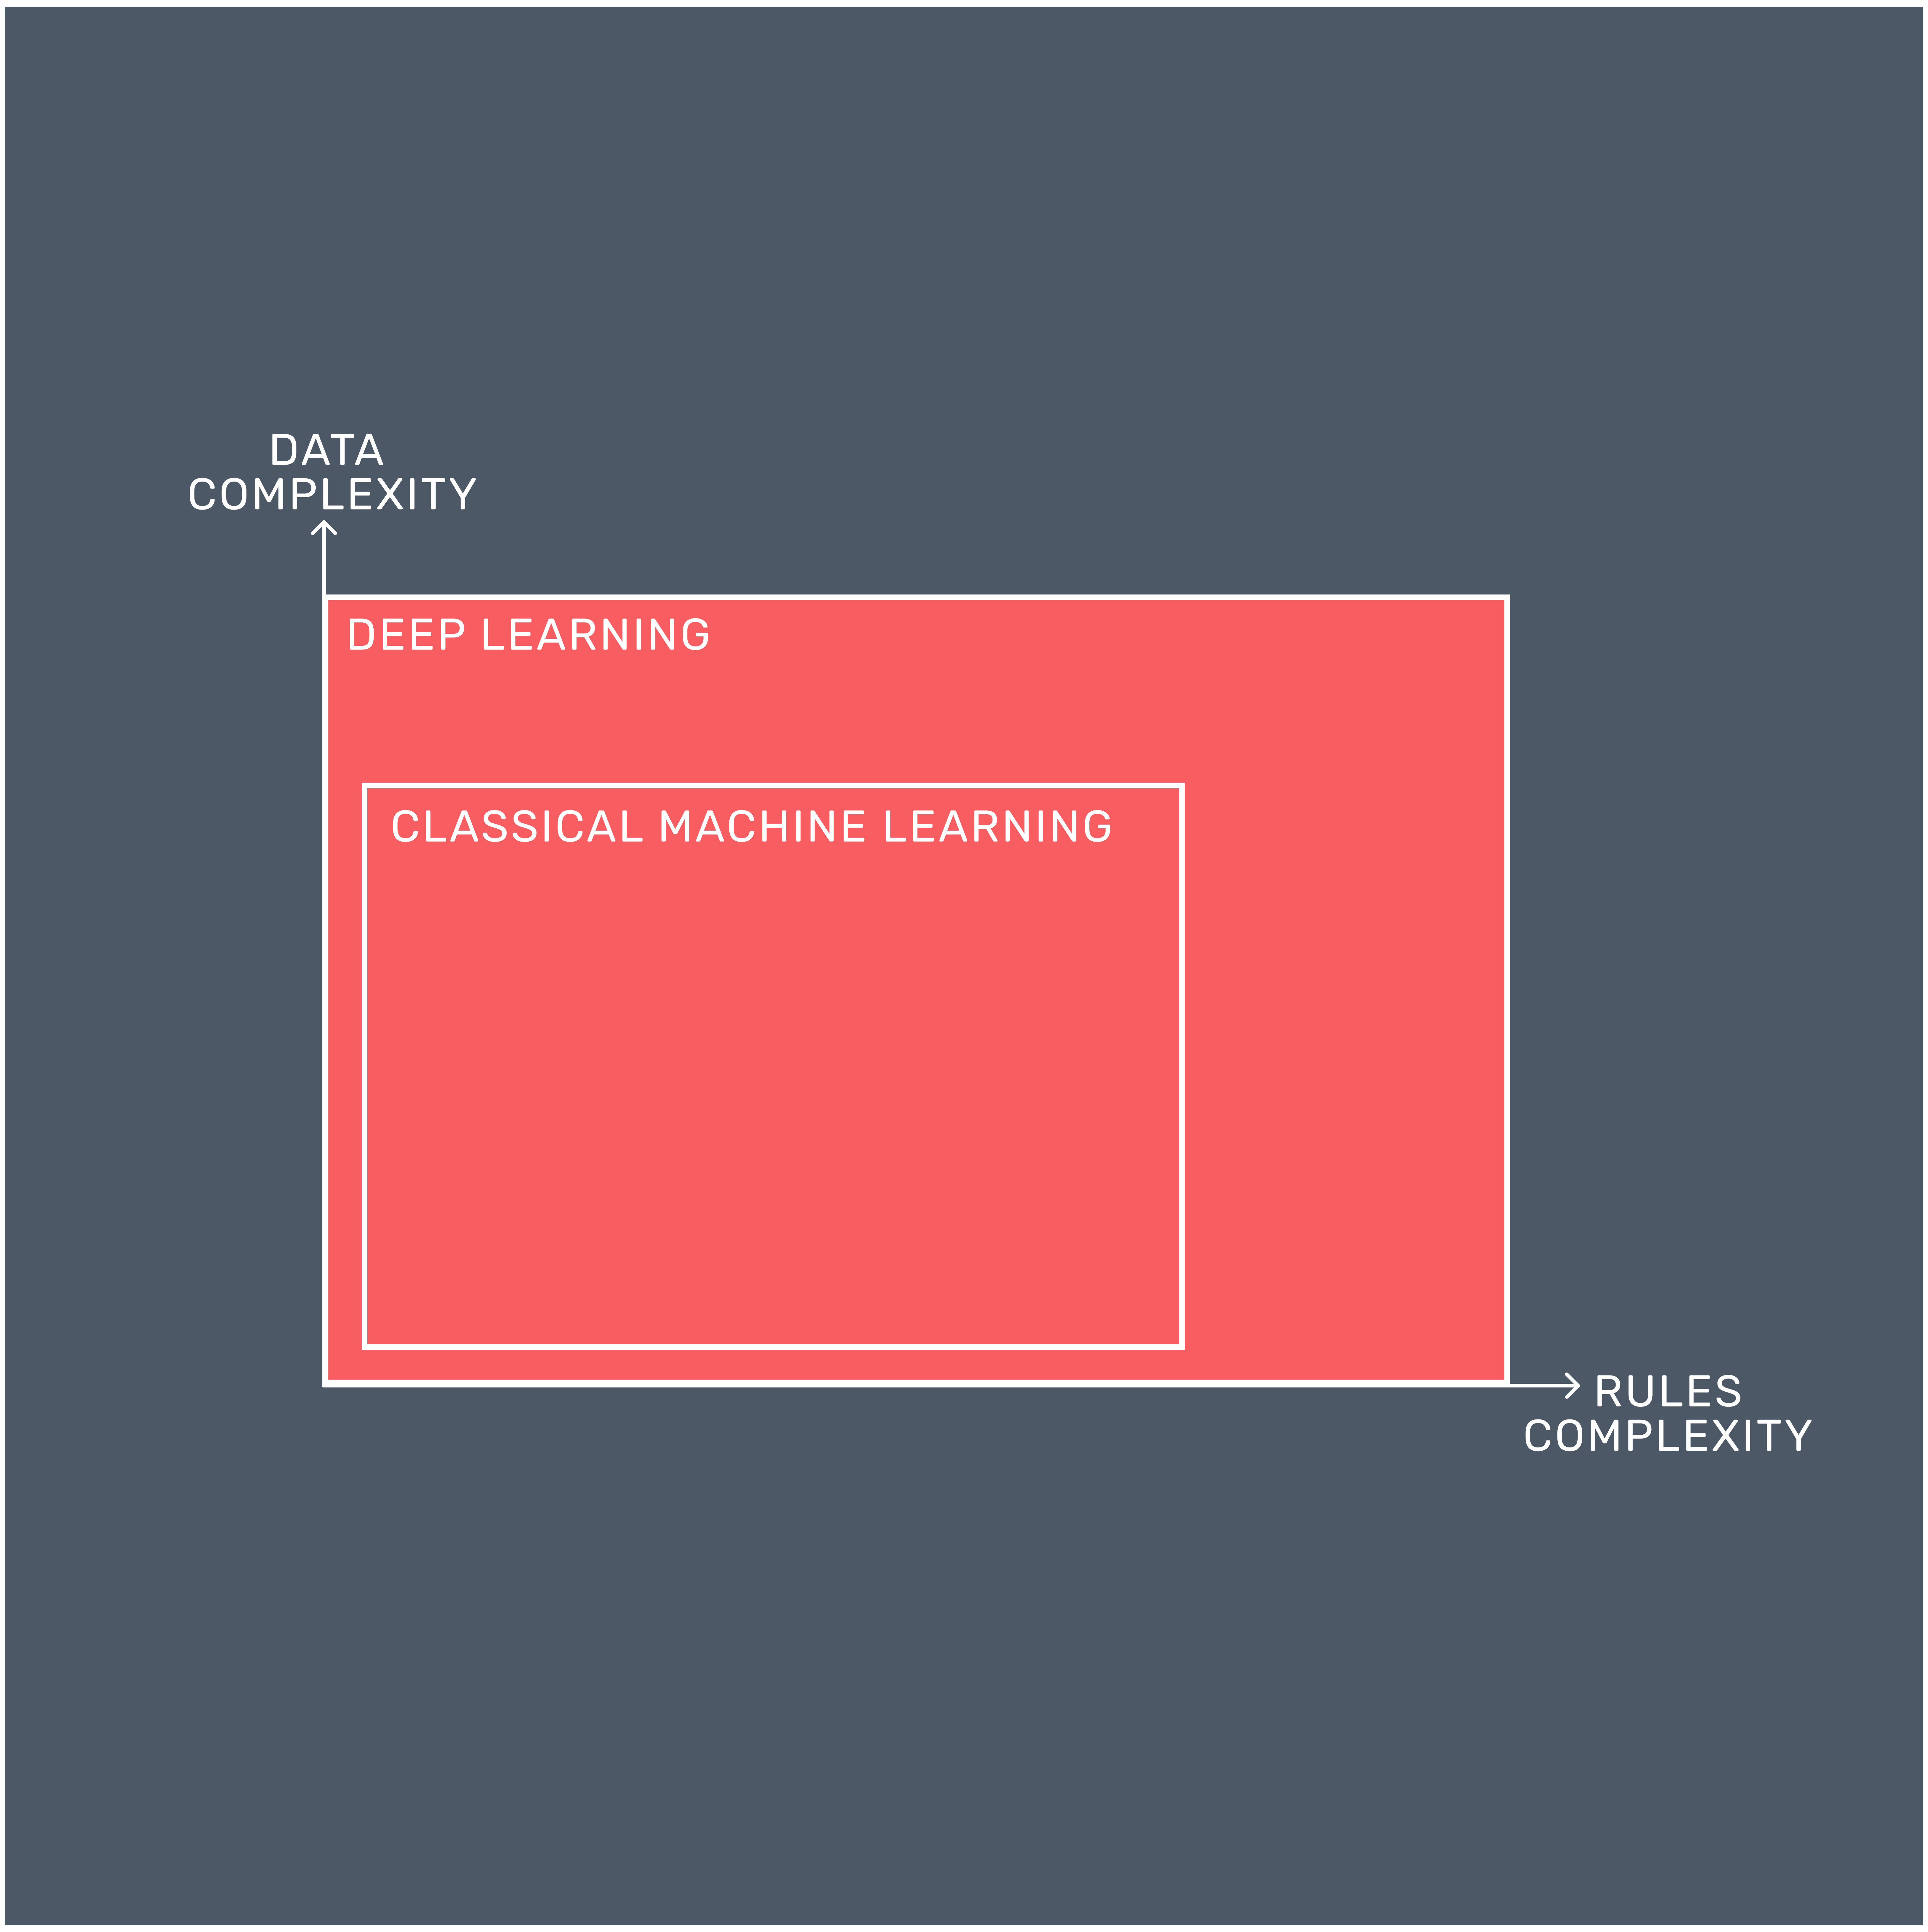

WHAT IS DEEP LEARNING?

Within machine learning, there are various types of algorithms. Think of machine learning algorithms as competing techniques to get the best out of the data. Some algorithms are better in certain aspects, but there’s not one that's the best in all departments.

Deep learning is a type of algorithm that's adaptable to varying complexities of data and rules.

It’s not necessarily the most accurate, but it's extremely adaptable. And this comes from its modular and flexible form, which will become evident throughout this book.

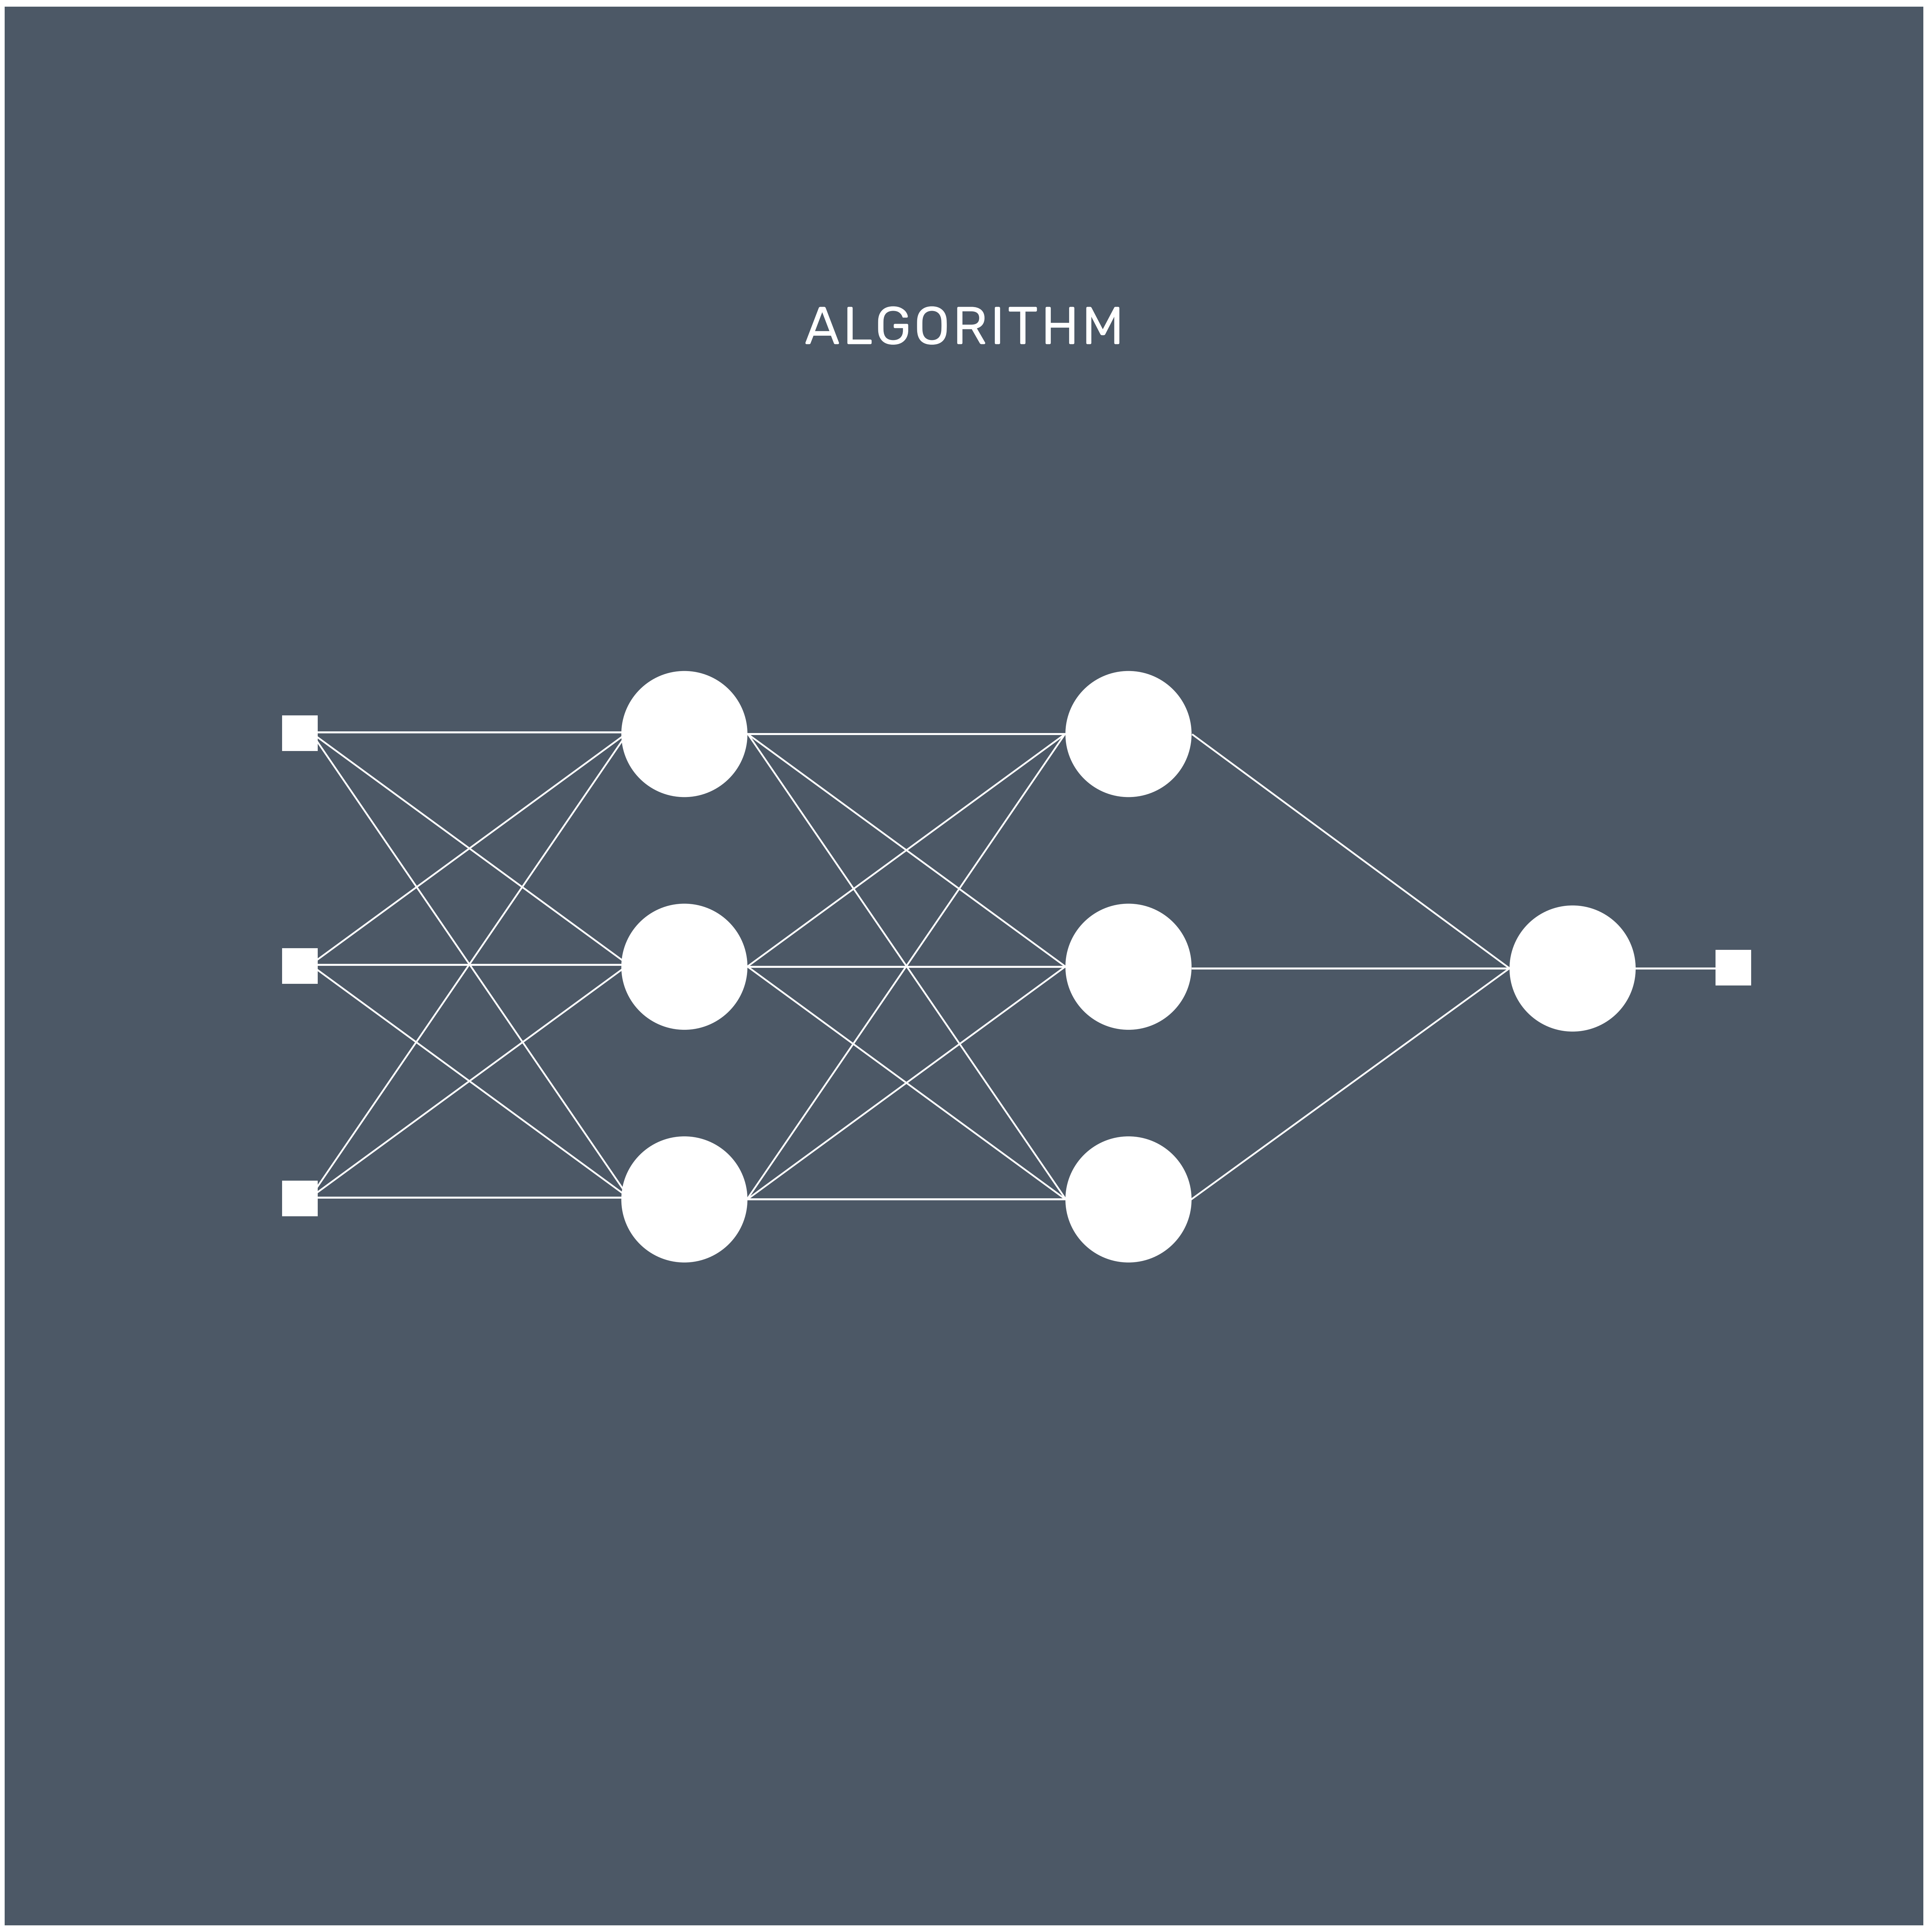

ALGORITHM

In fact, deep learning has revived the push toward Artificial Intelligence (AI) over the past decade.

The progress is gathering pace now is because of three main reasons. The first is the algorithm, which in truth, has been around for many decades.

But that alone is not enough.

DATA

The second reason is data.

The impetus came from the Internet and followed by social media, smartphones, digital transformation, and a long list of other waves of innovation. They produce new forms of data that we've never seen before, generated in large volumes.

This data contains invaluable information that we can now extract with the help of algorithms.

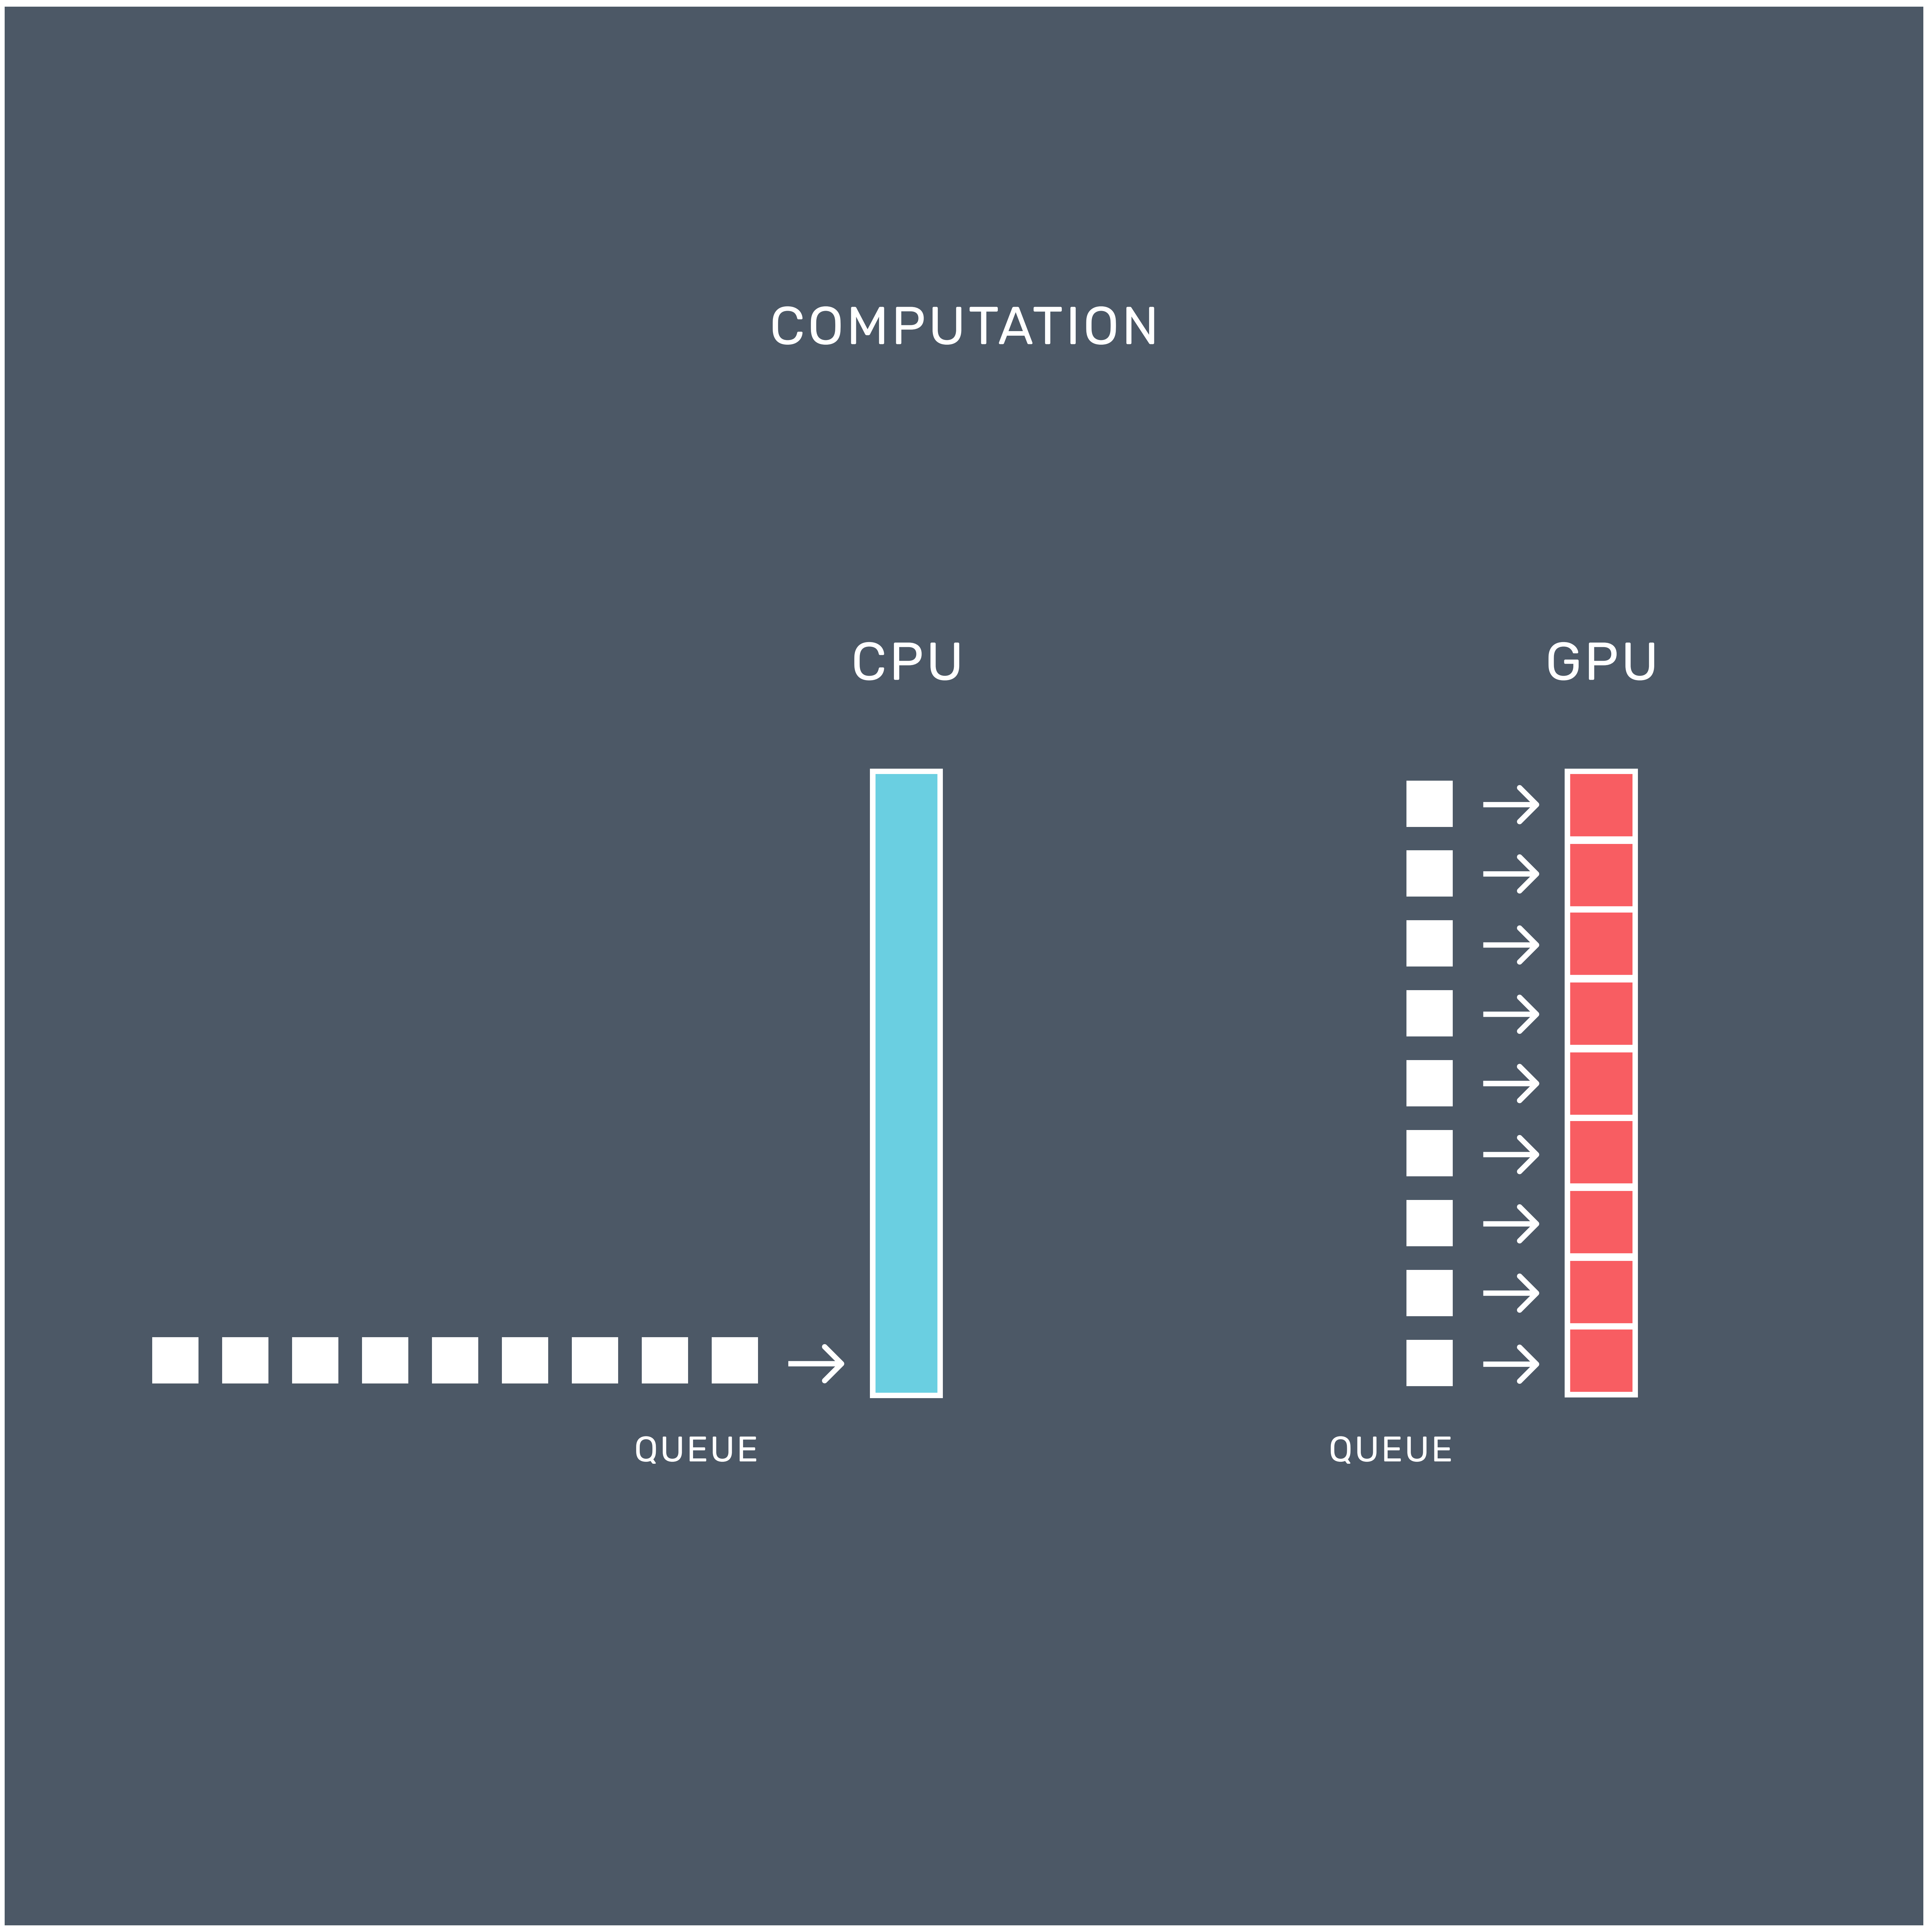

COMPUTATION

The third reason is computational power.

Machine learning involves performing a significant amount of mathematical computation on the data. In deep learning, this is multiplied many times over. The standard Central Processing Unit (CPU) architecture is not capable of handling this task efficiently.

Enter the Graphics Processing Units (GPU). Originally designed for games, it has emerged as the perfect solution for deep learning.

This is a hot area of research as we speak. Even more efficient hardware designs are yet to come.

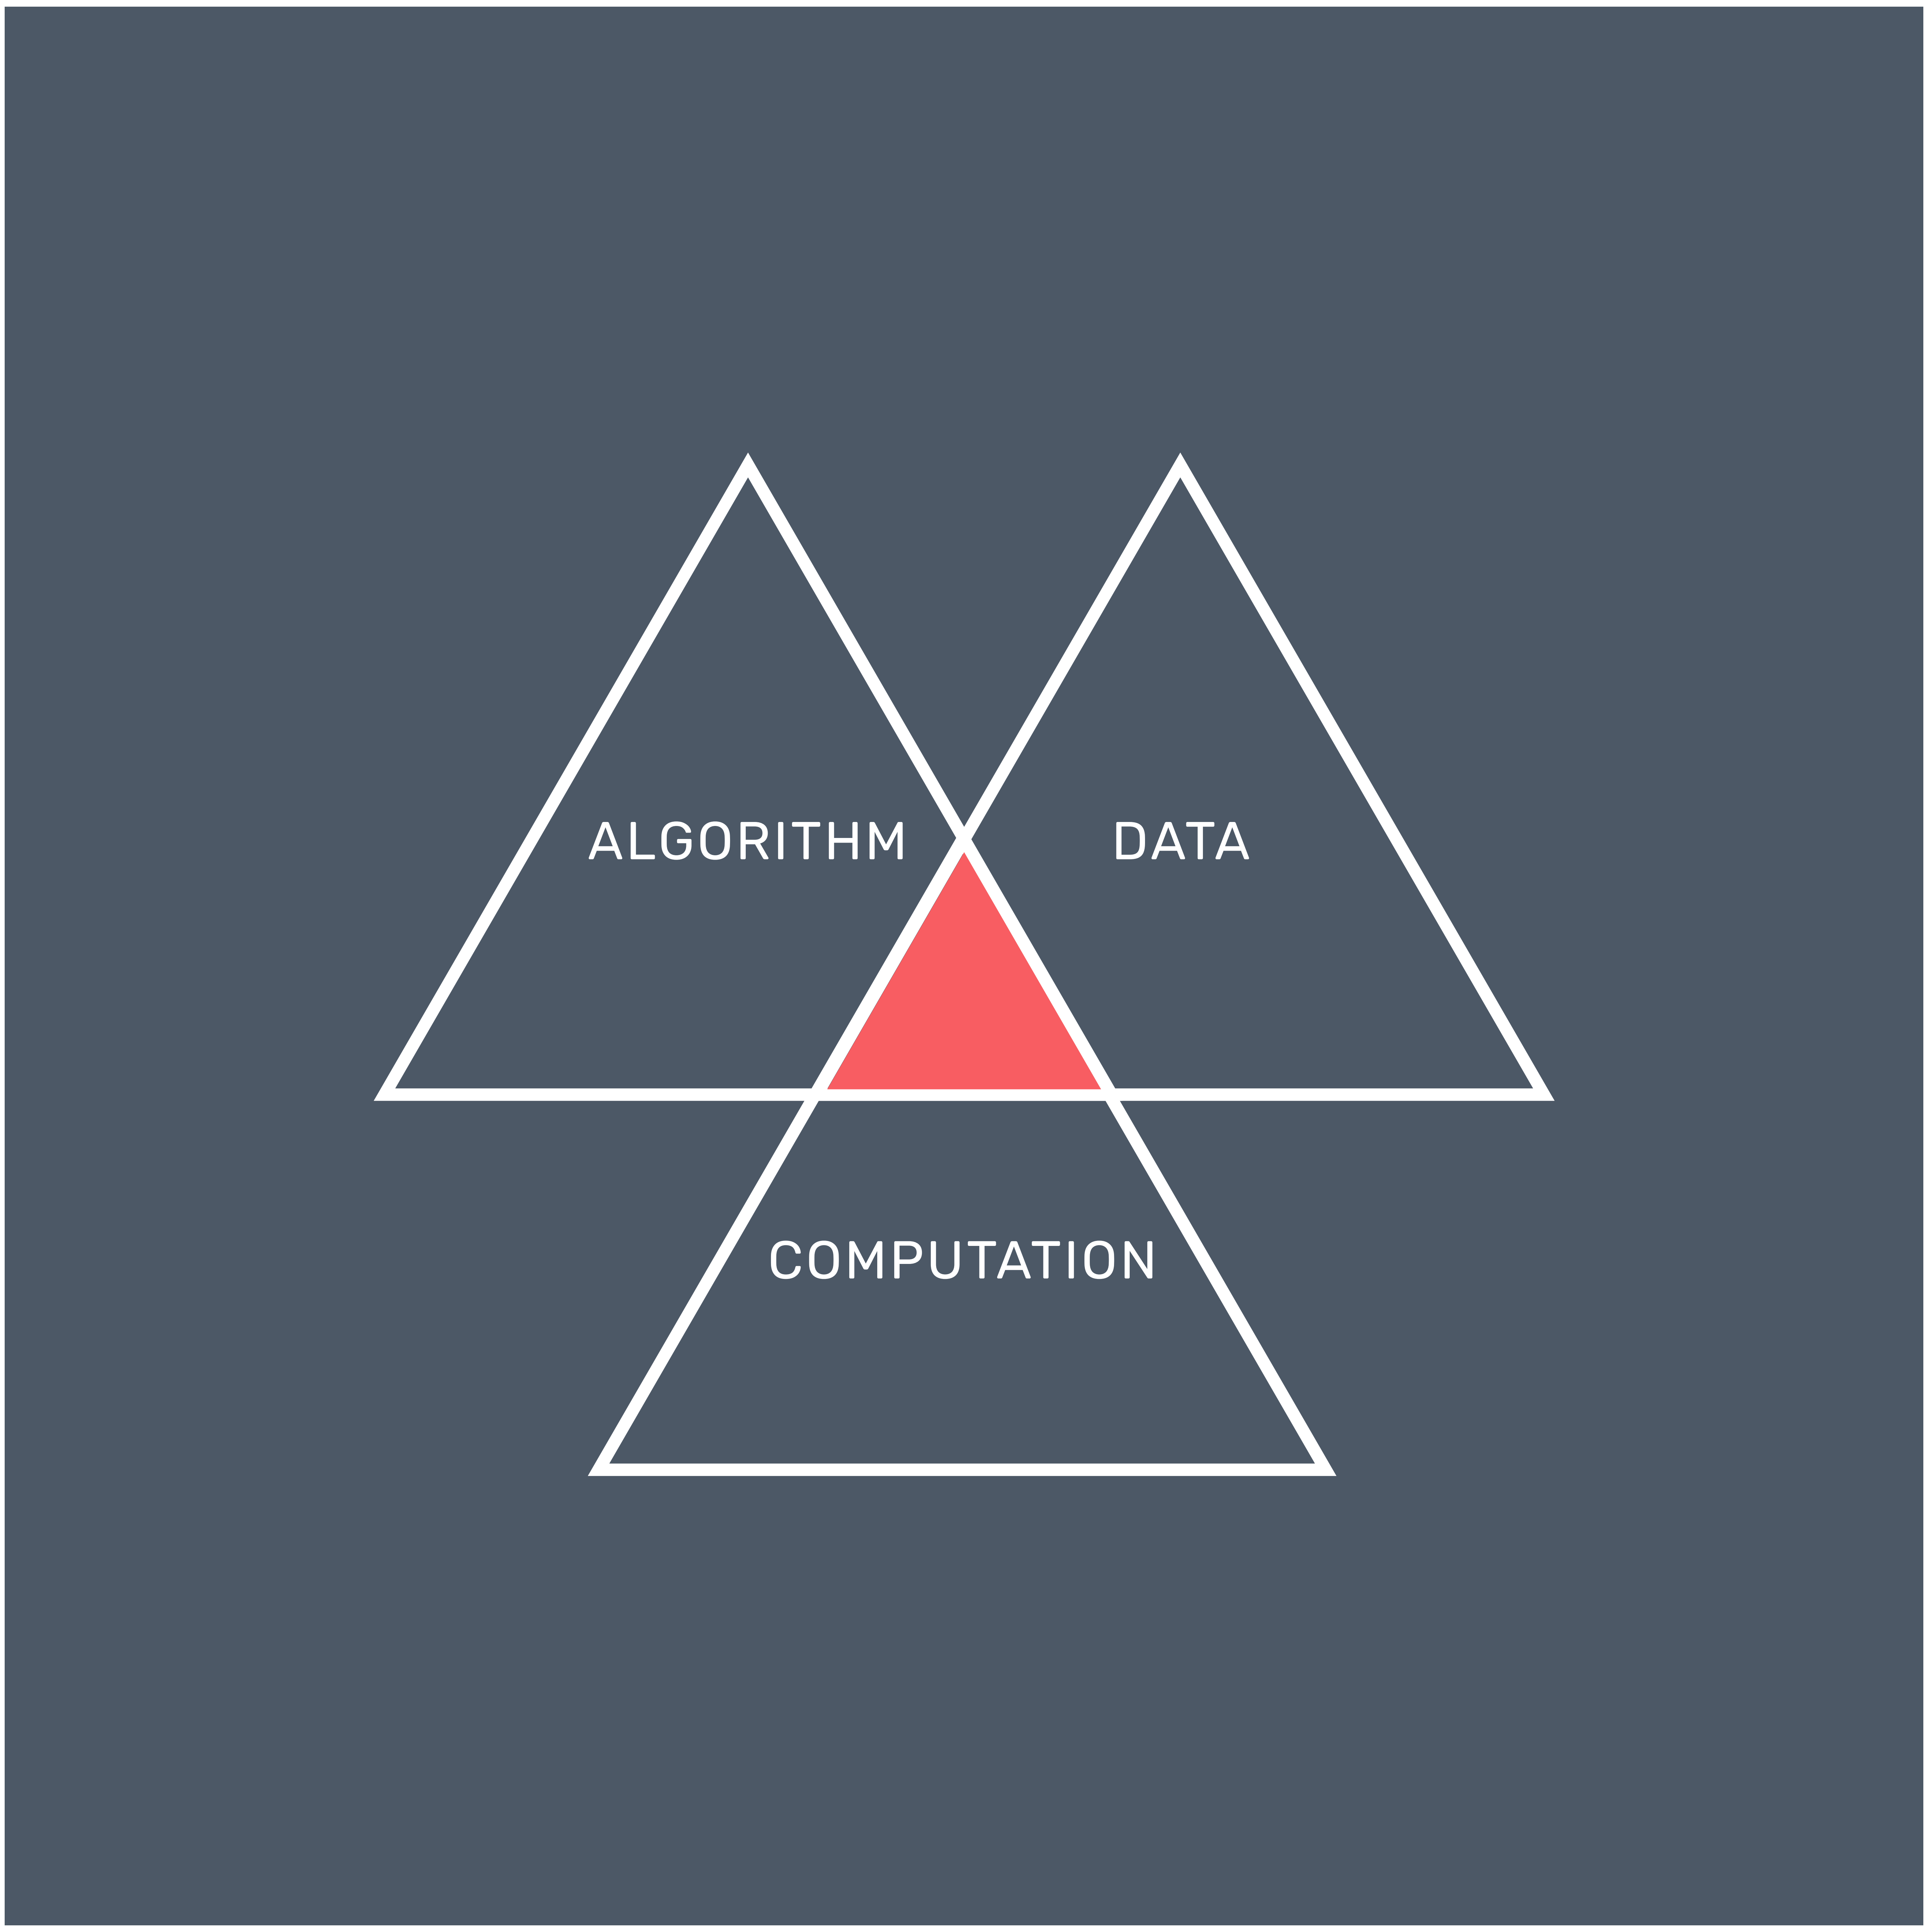

THE DRIVING FORCES

Together, these three factors are the driving forces behind today's rapid advances in deep learning.

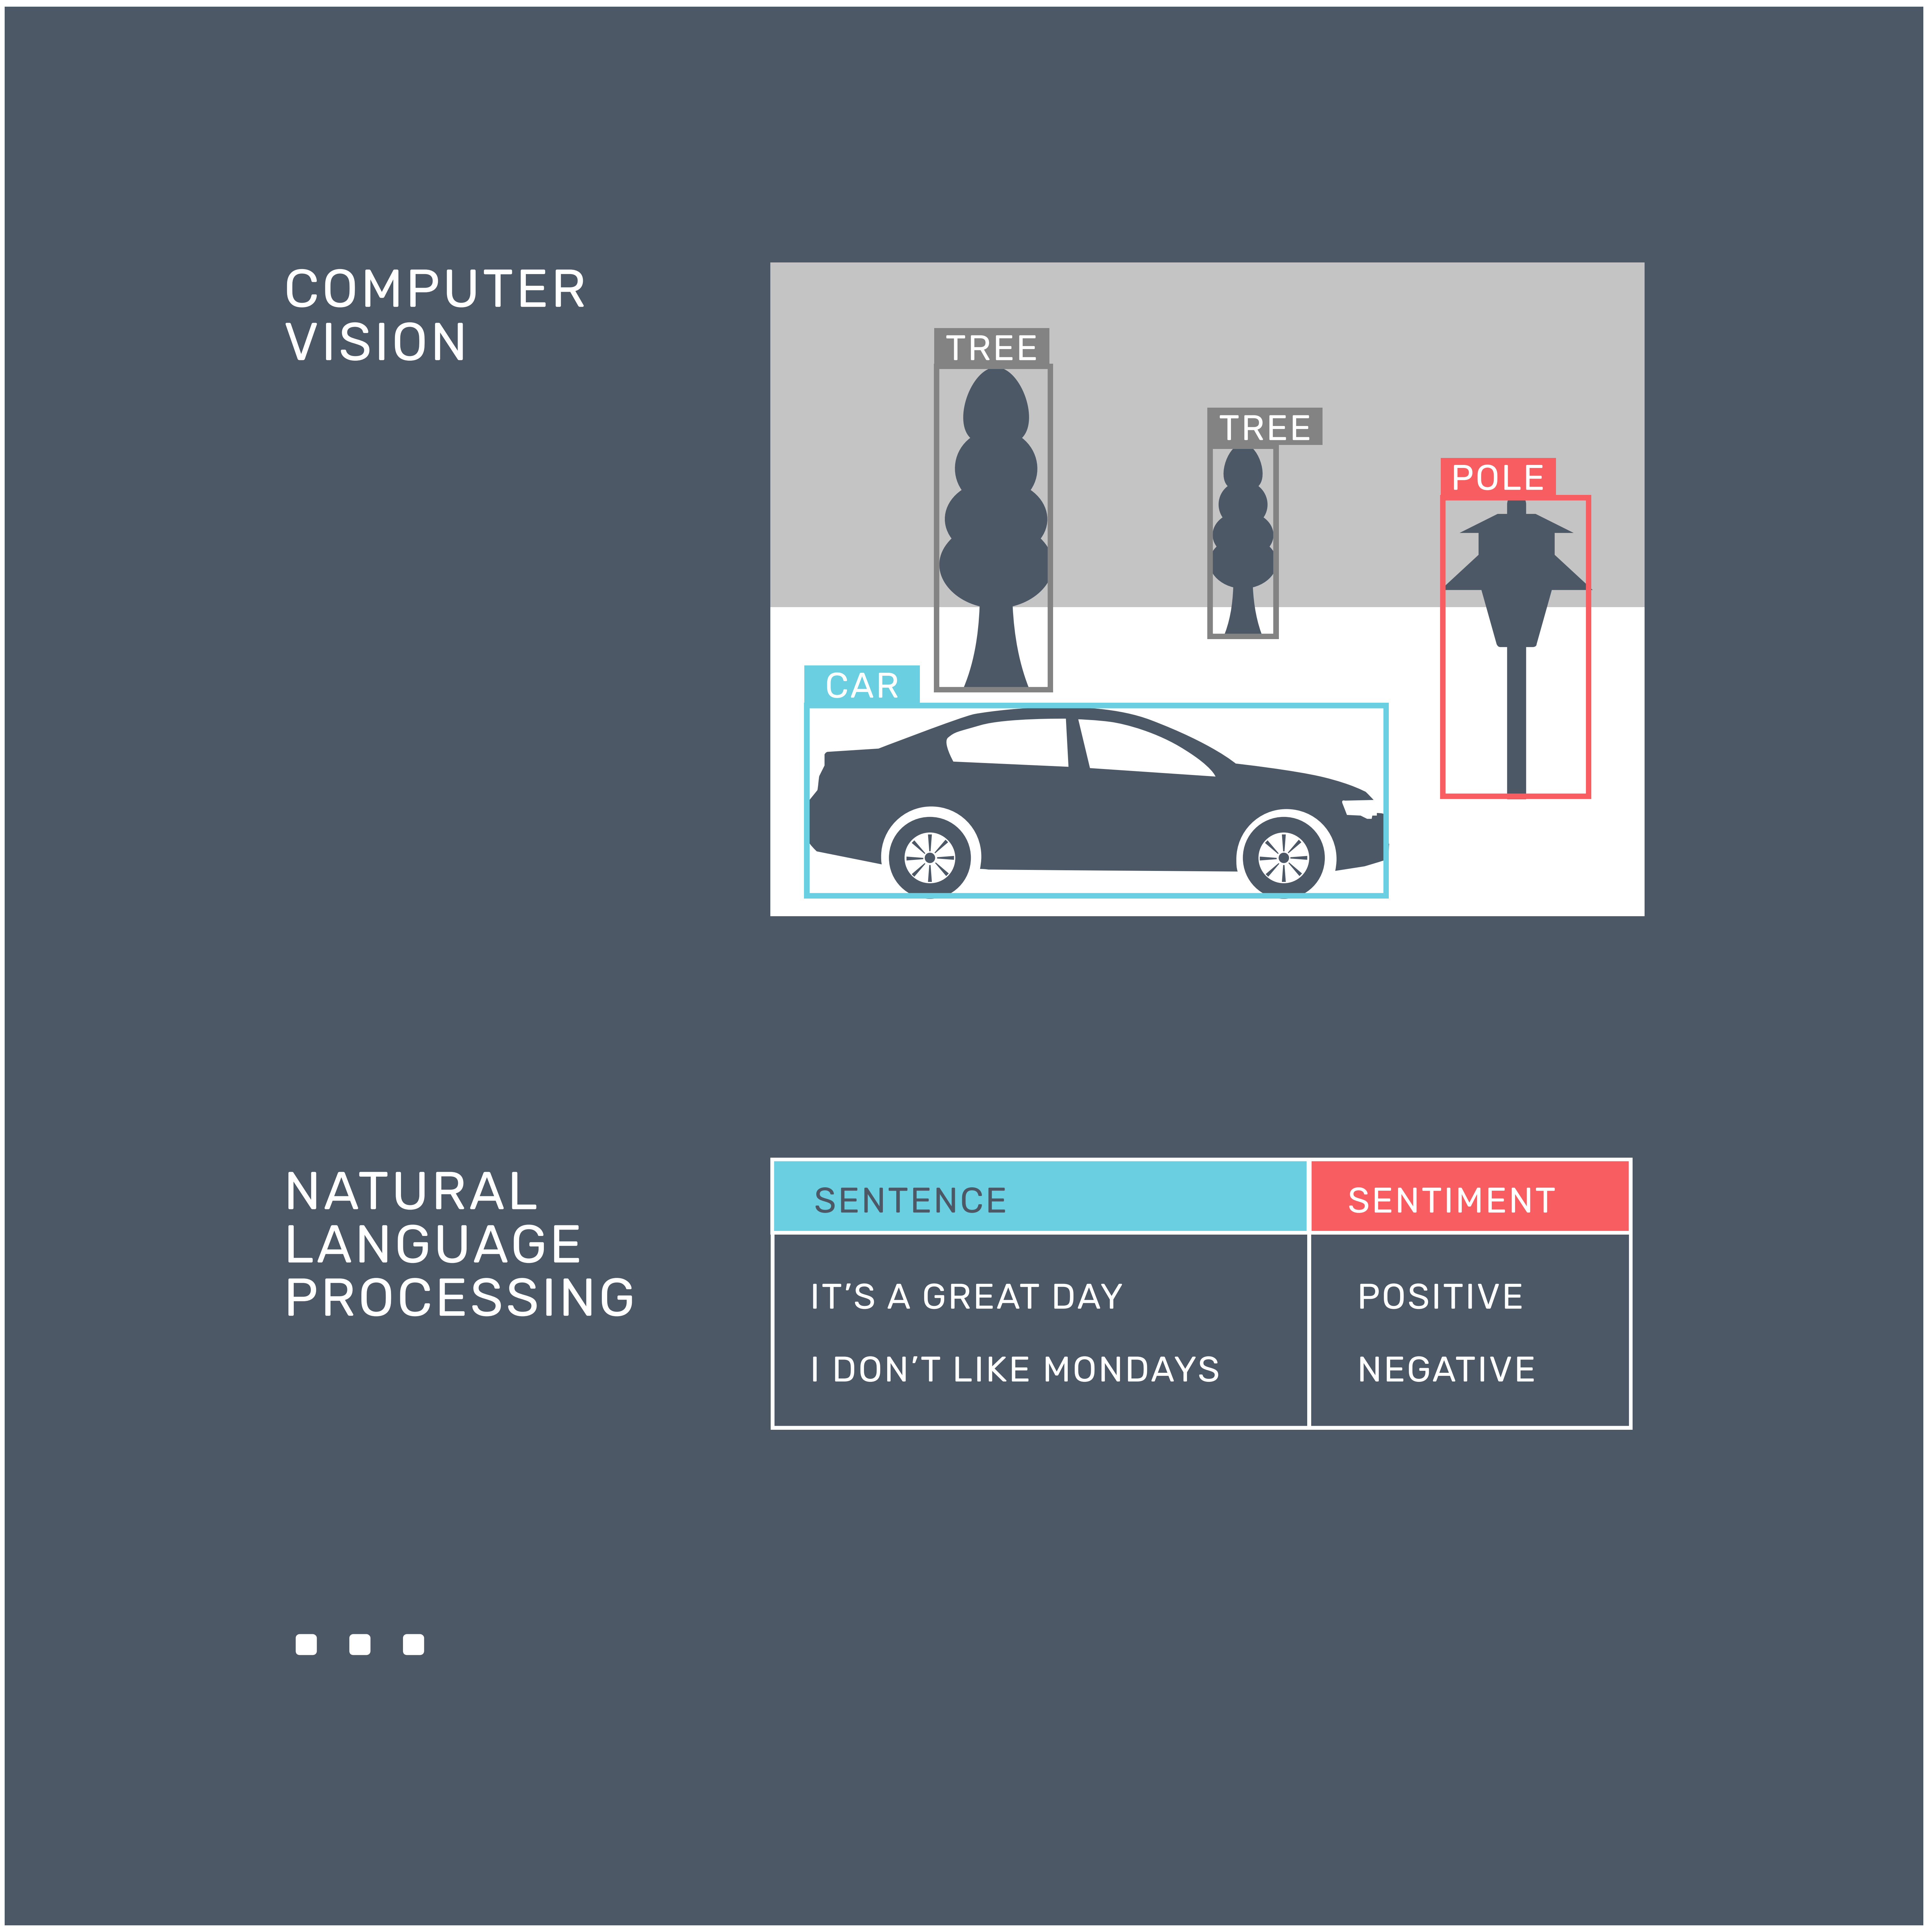

APPLICATIONS

Today, there are widespread applications in computer vision, natural language processing, business automation, and beyond. And it is just the beginning.

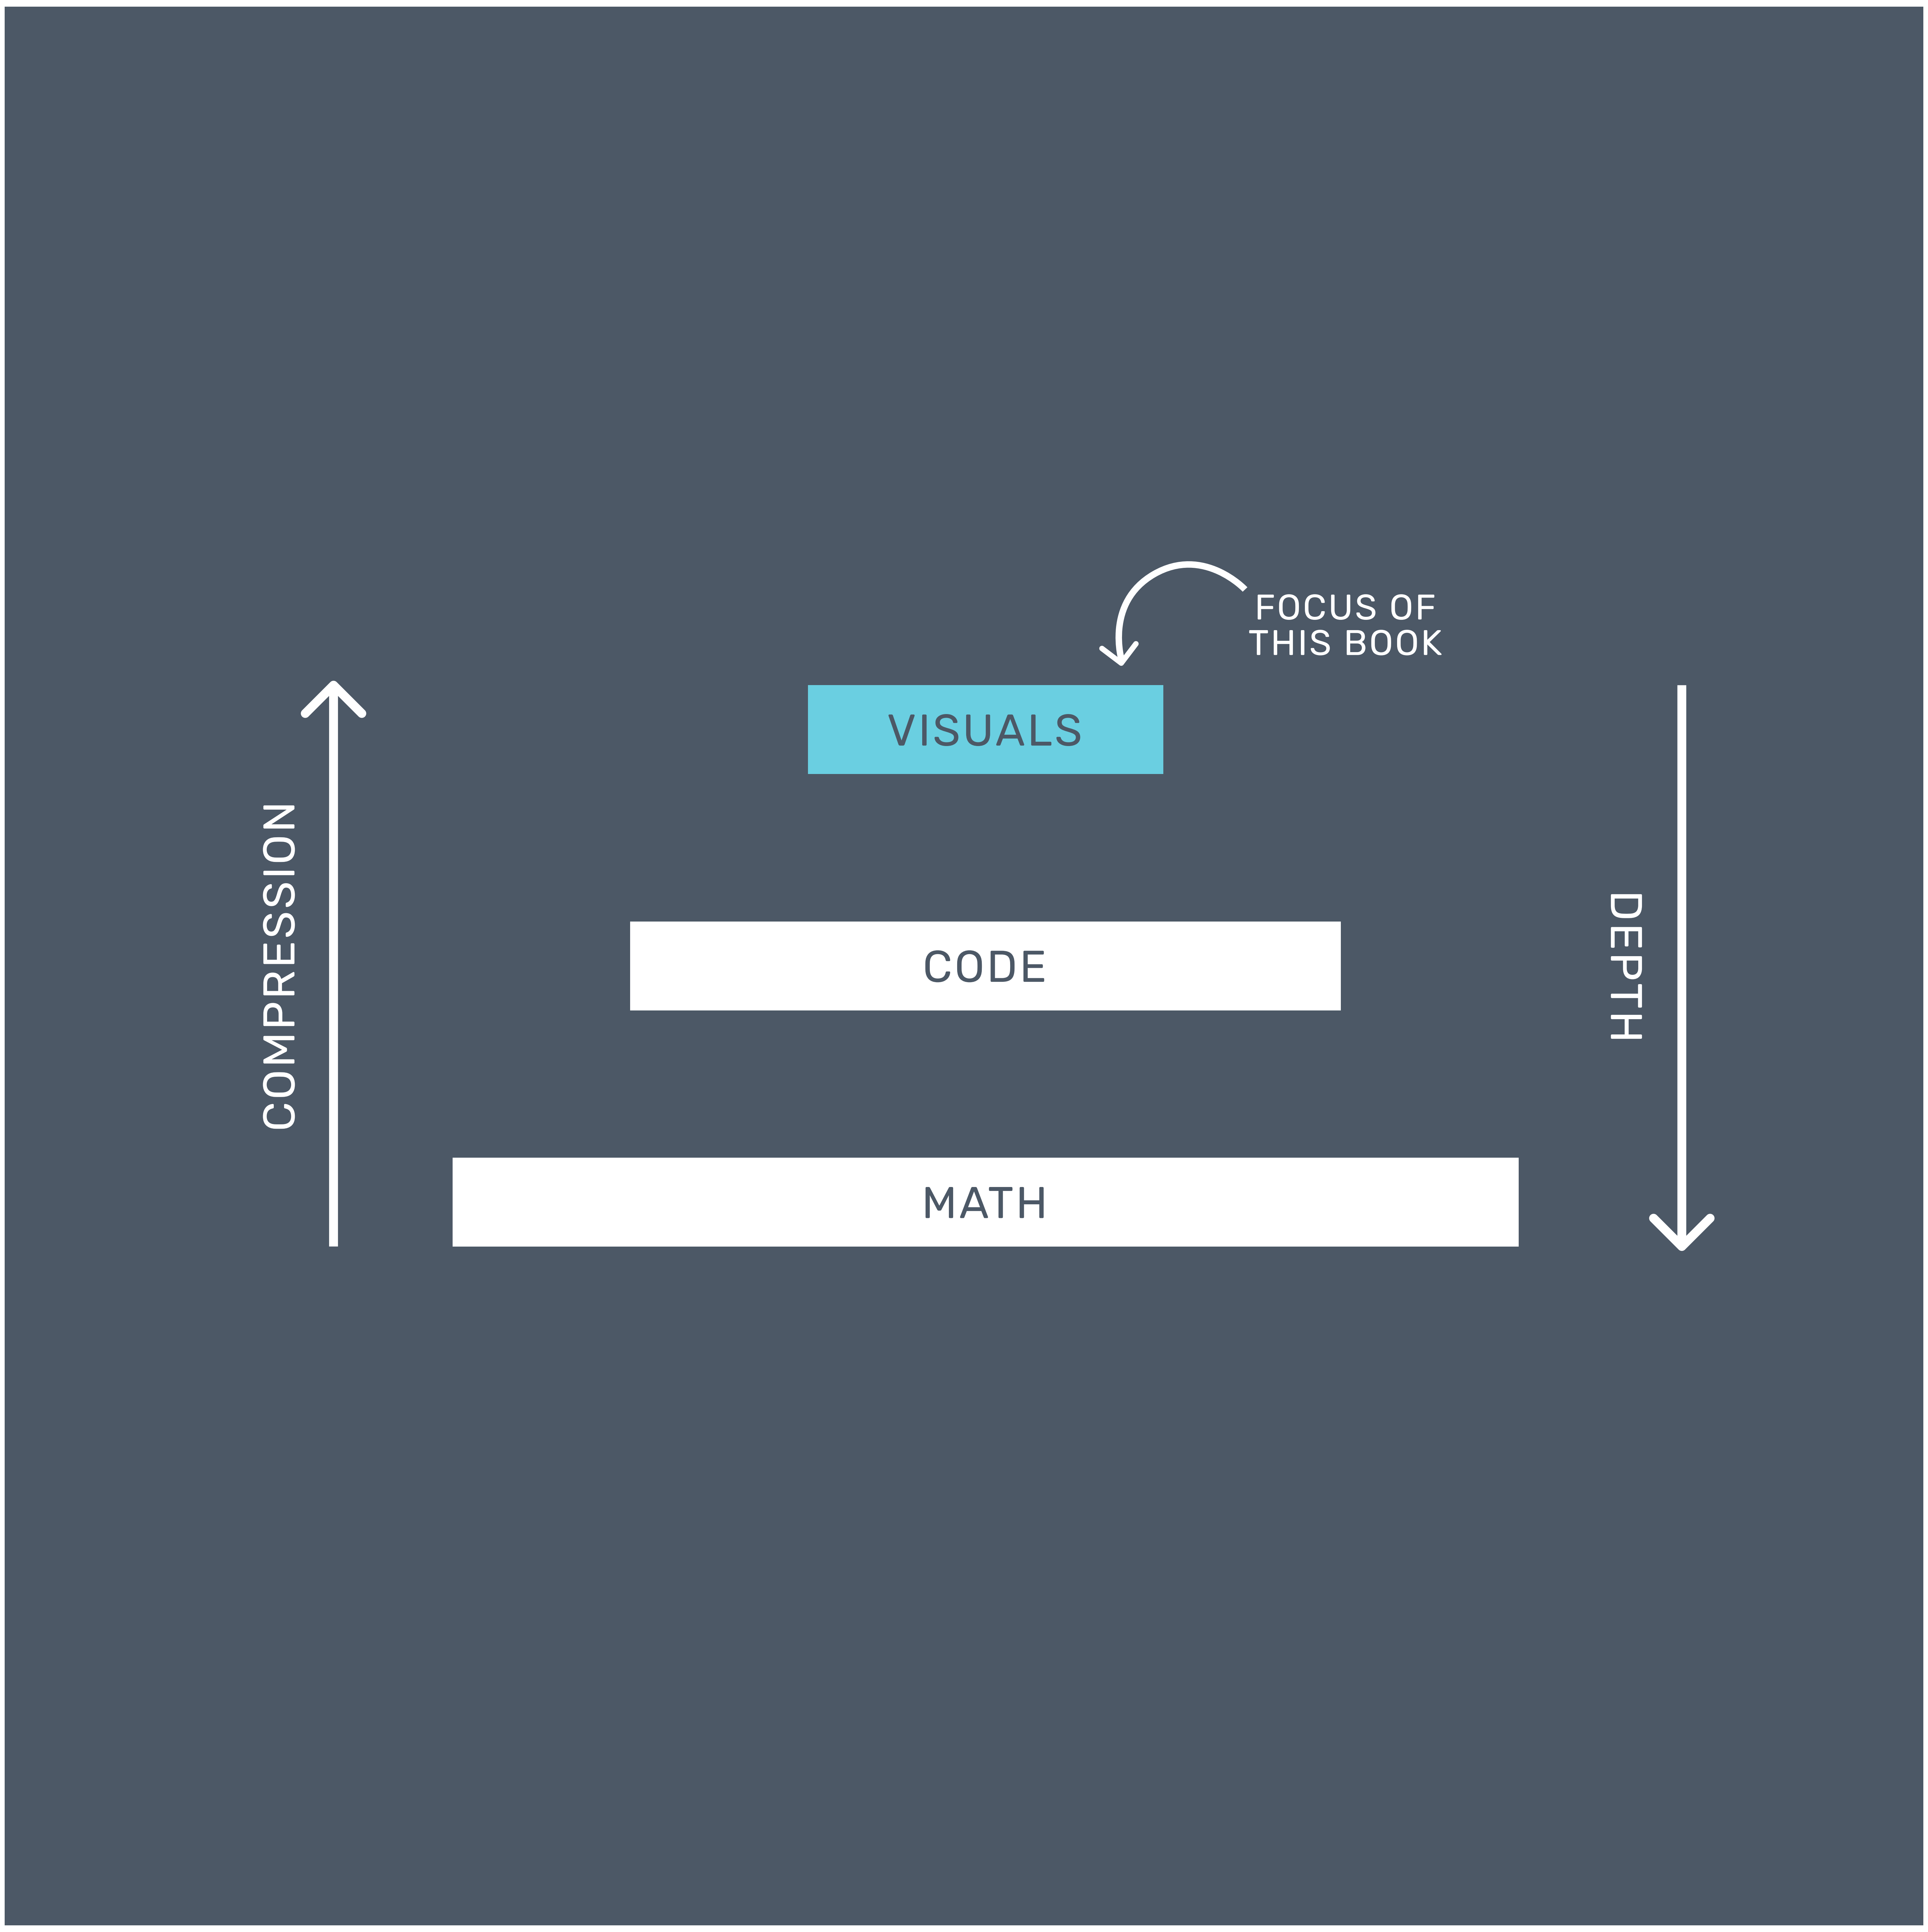

WHAT CAN YOU EXPECT FROM THIS BOOK?

By the end of this book, you will be able to build a visual intuition about deep learning and neural networks.

This book doesn't cover mathematical proofs and code examples. As you advance your learning further, these are the domains you should progress into. They will provide you with the depth you need to be successful in this field.

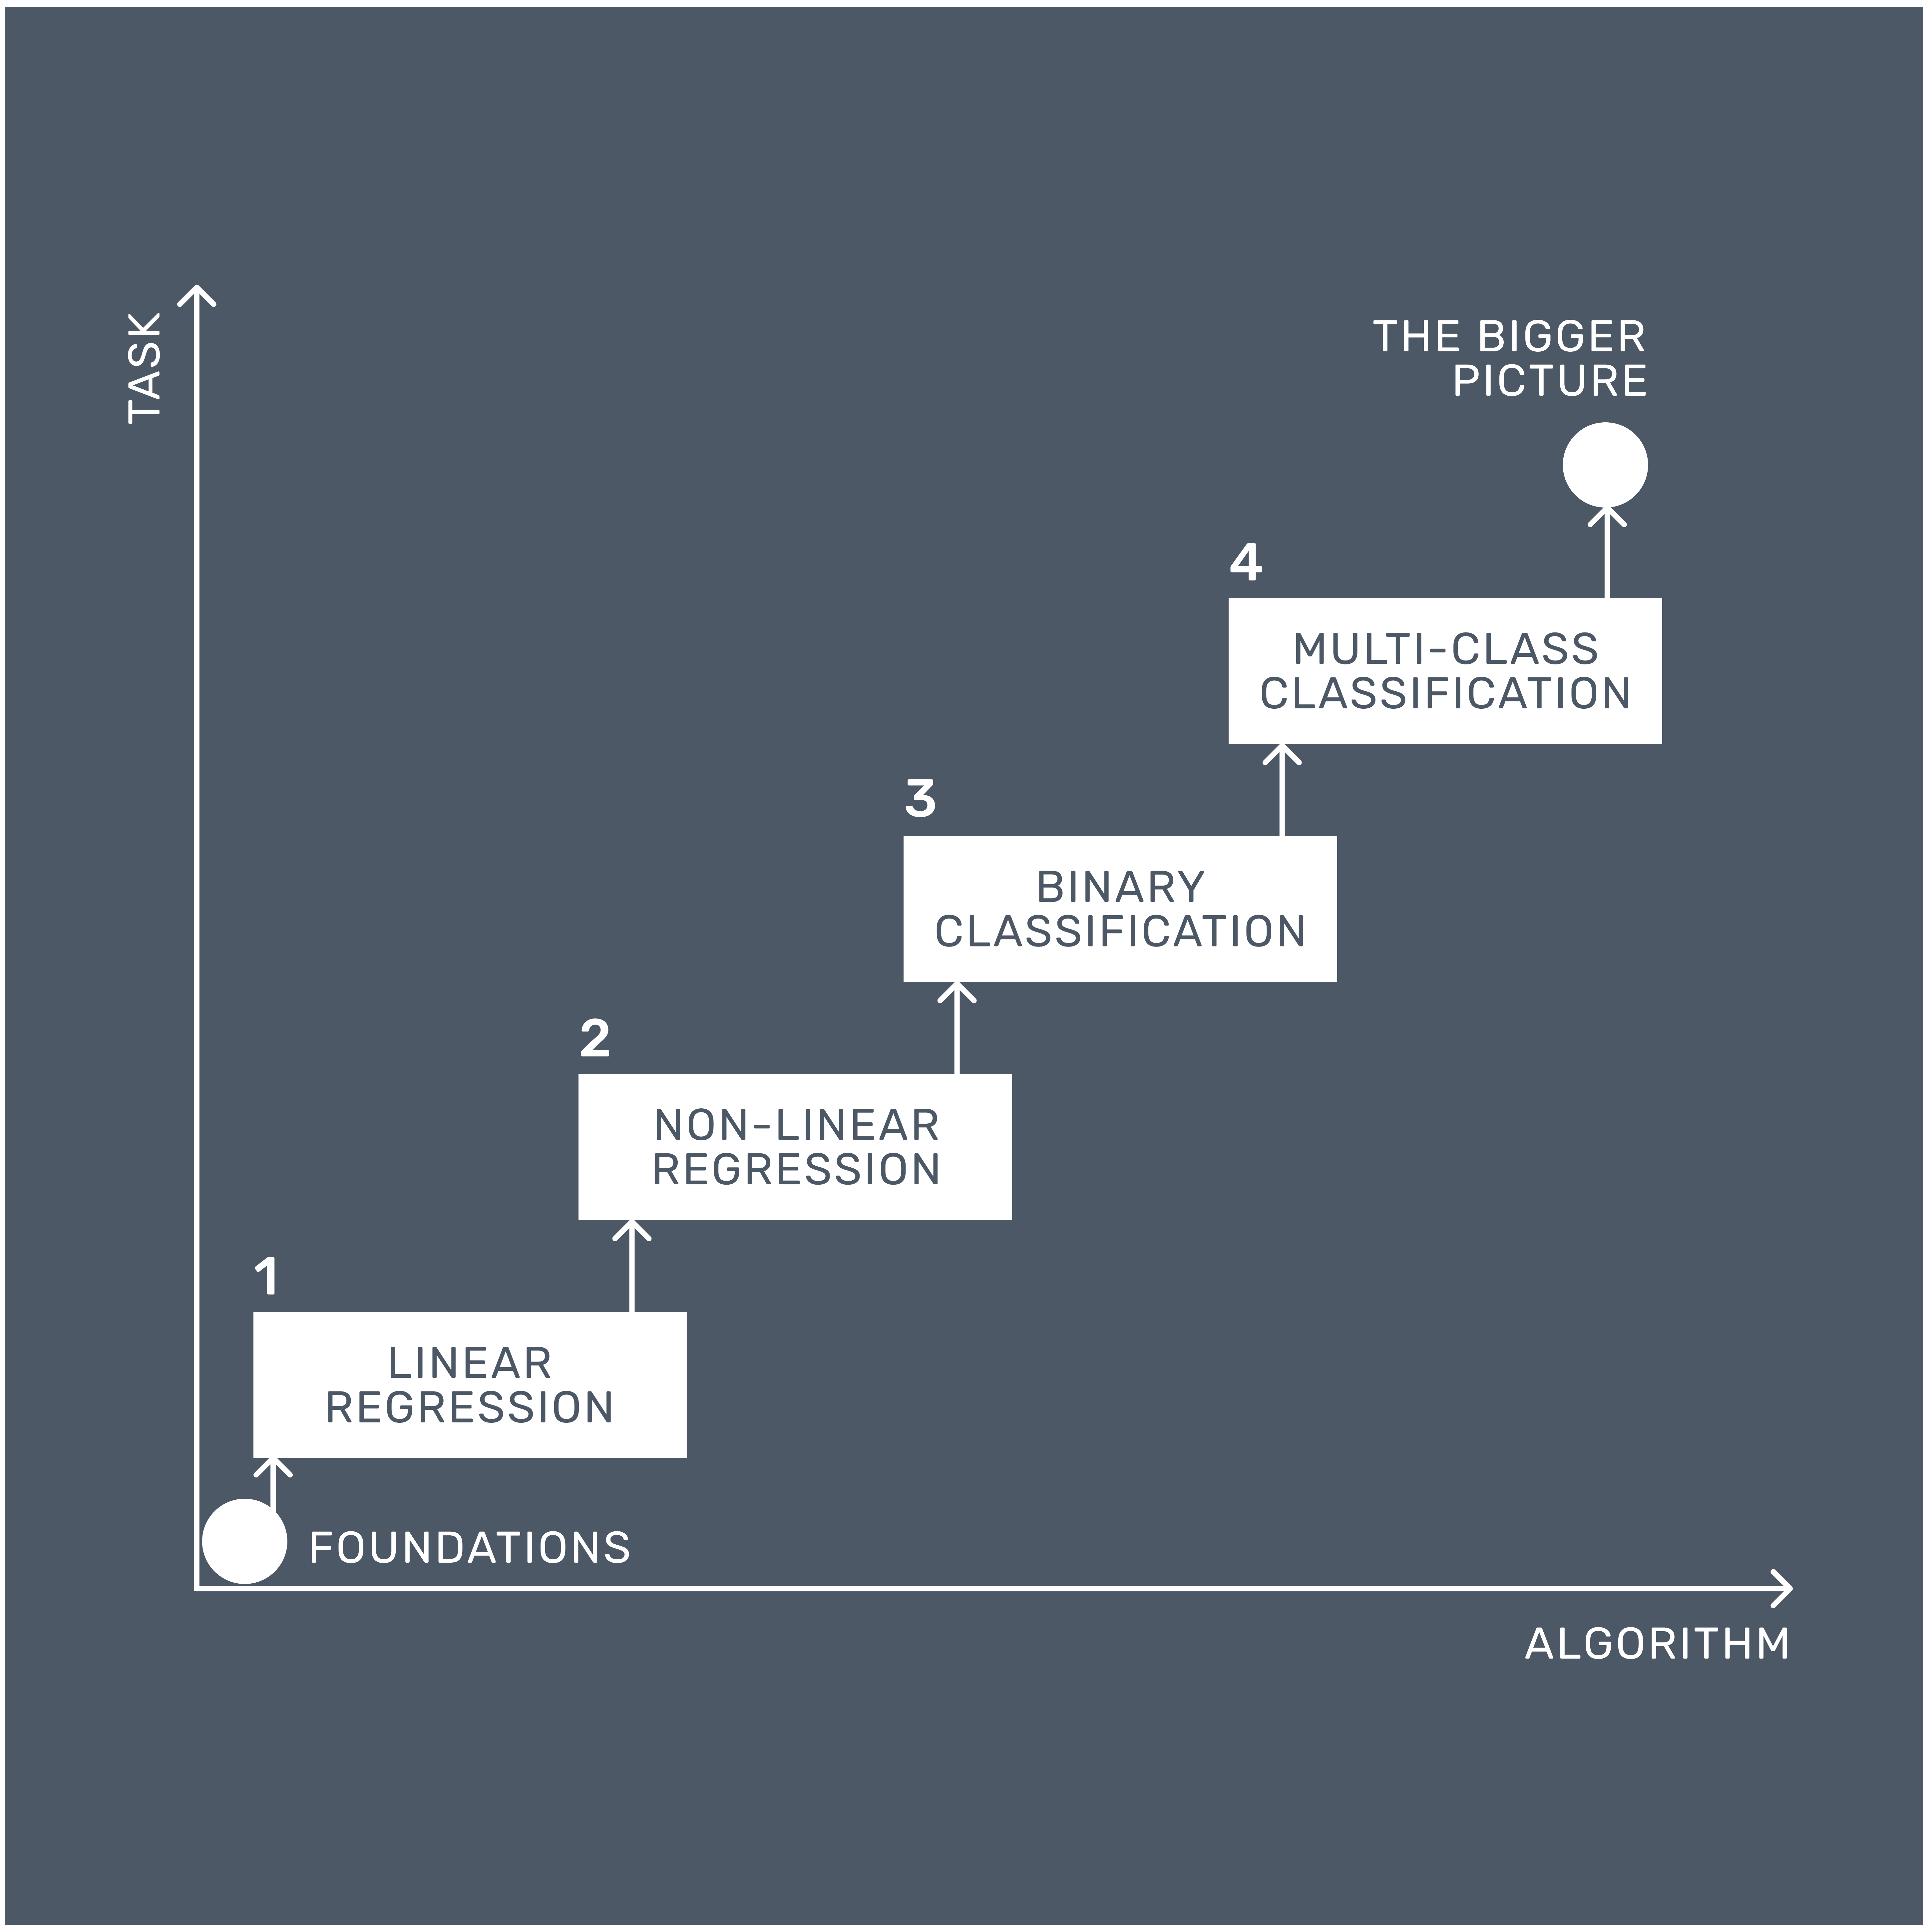

ROADMAP

We'll see how deep learning works via four tasks - linear regression, non-linear regression, binary classification, multi-class classification.

They are correspondingly split into four chapters, in which new concepts are introduced one at a time and built upon the previous ones. Therefore, it is recommended that you read the chapters in sequence.

On either side of these four chapters, we'll have a short section for foundations and a final section where we take a brief look beyond those covered in the four chapters.

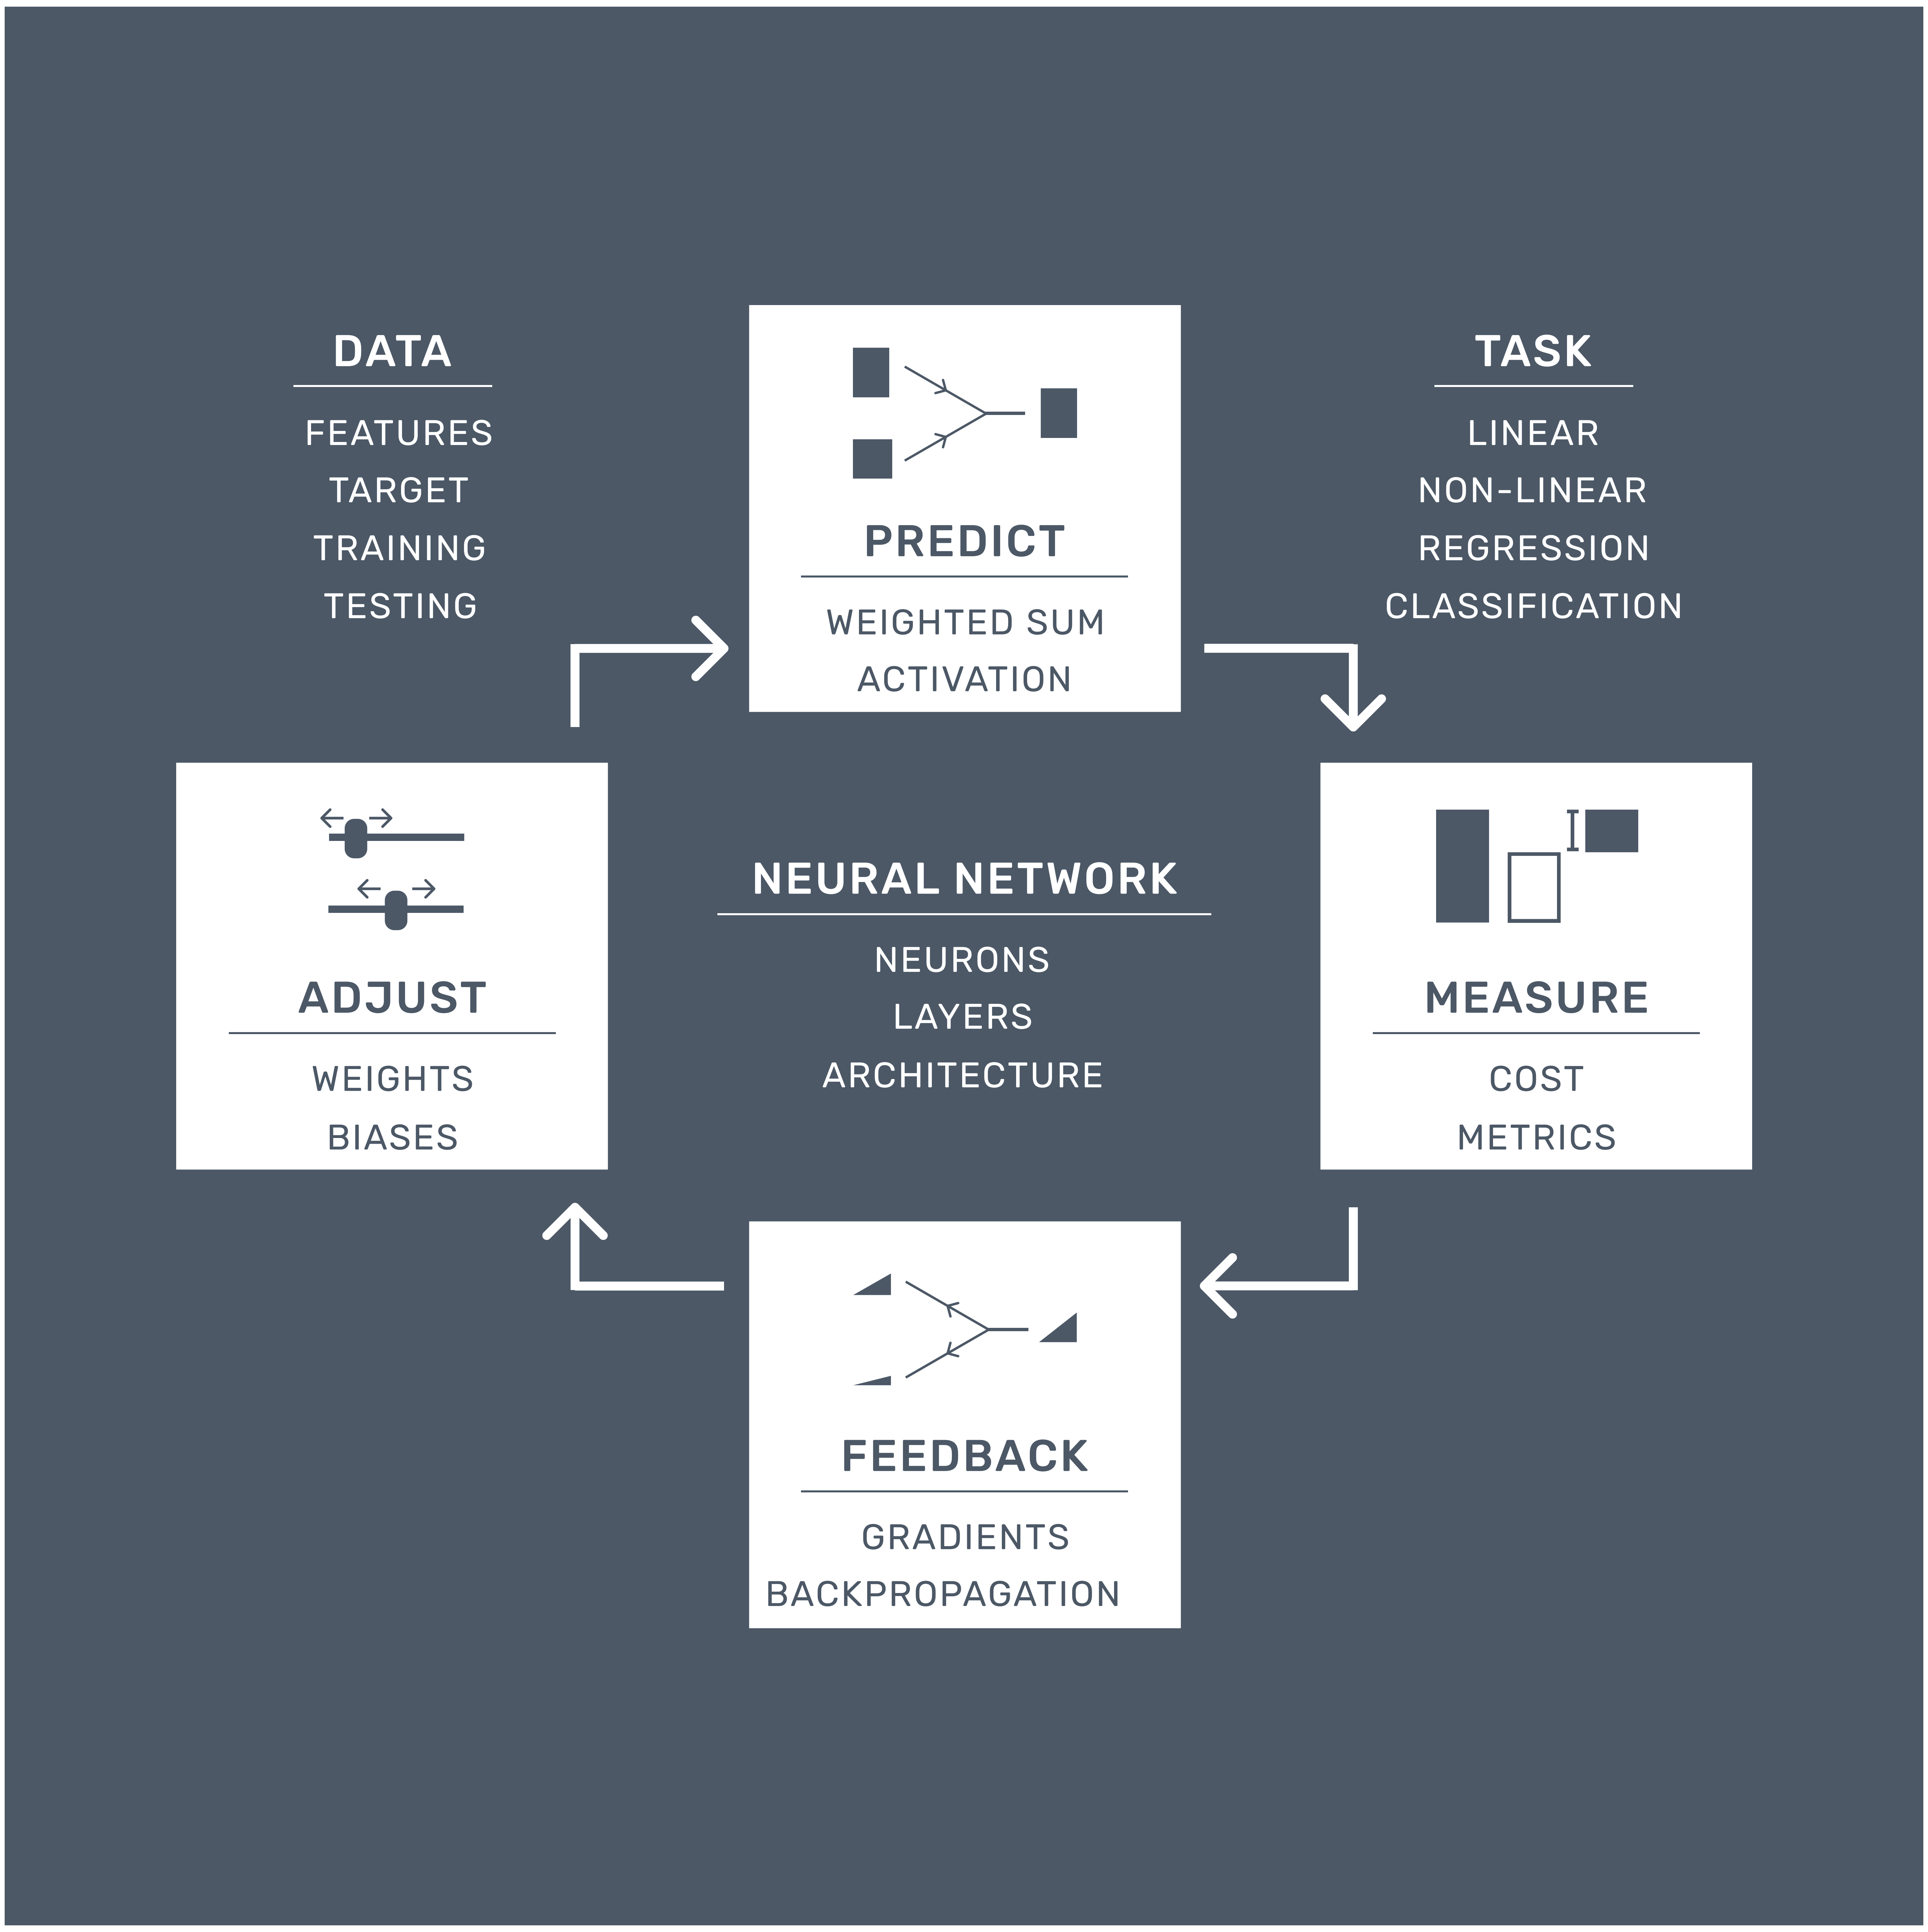

KEY CONCEPTS

Here is a summary of the key concepts that we'll explore in this book. As you go through the book, it'll be useful to return to this page from time to time to keep track of what you have learned.

Let's begin!

FOUNDATIONS

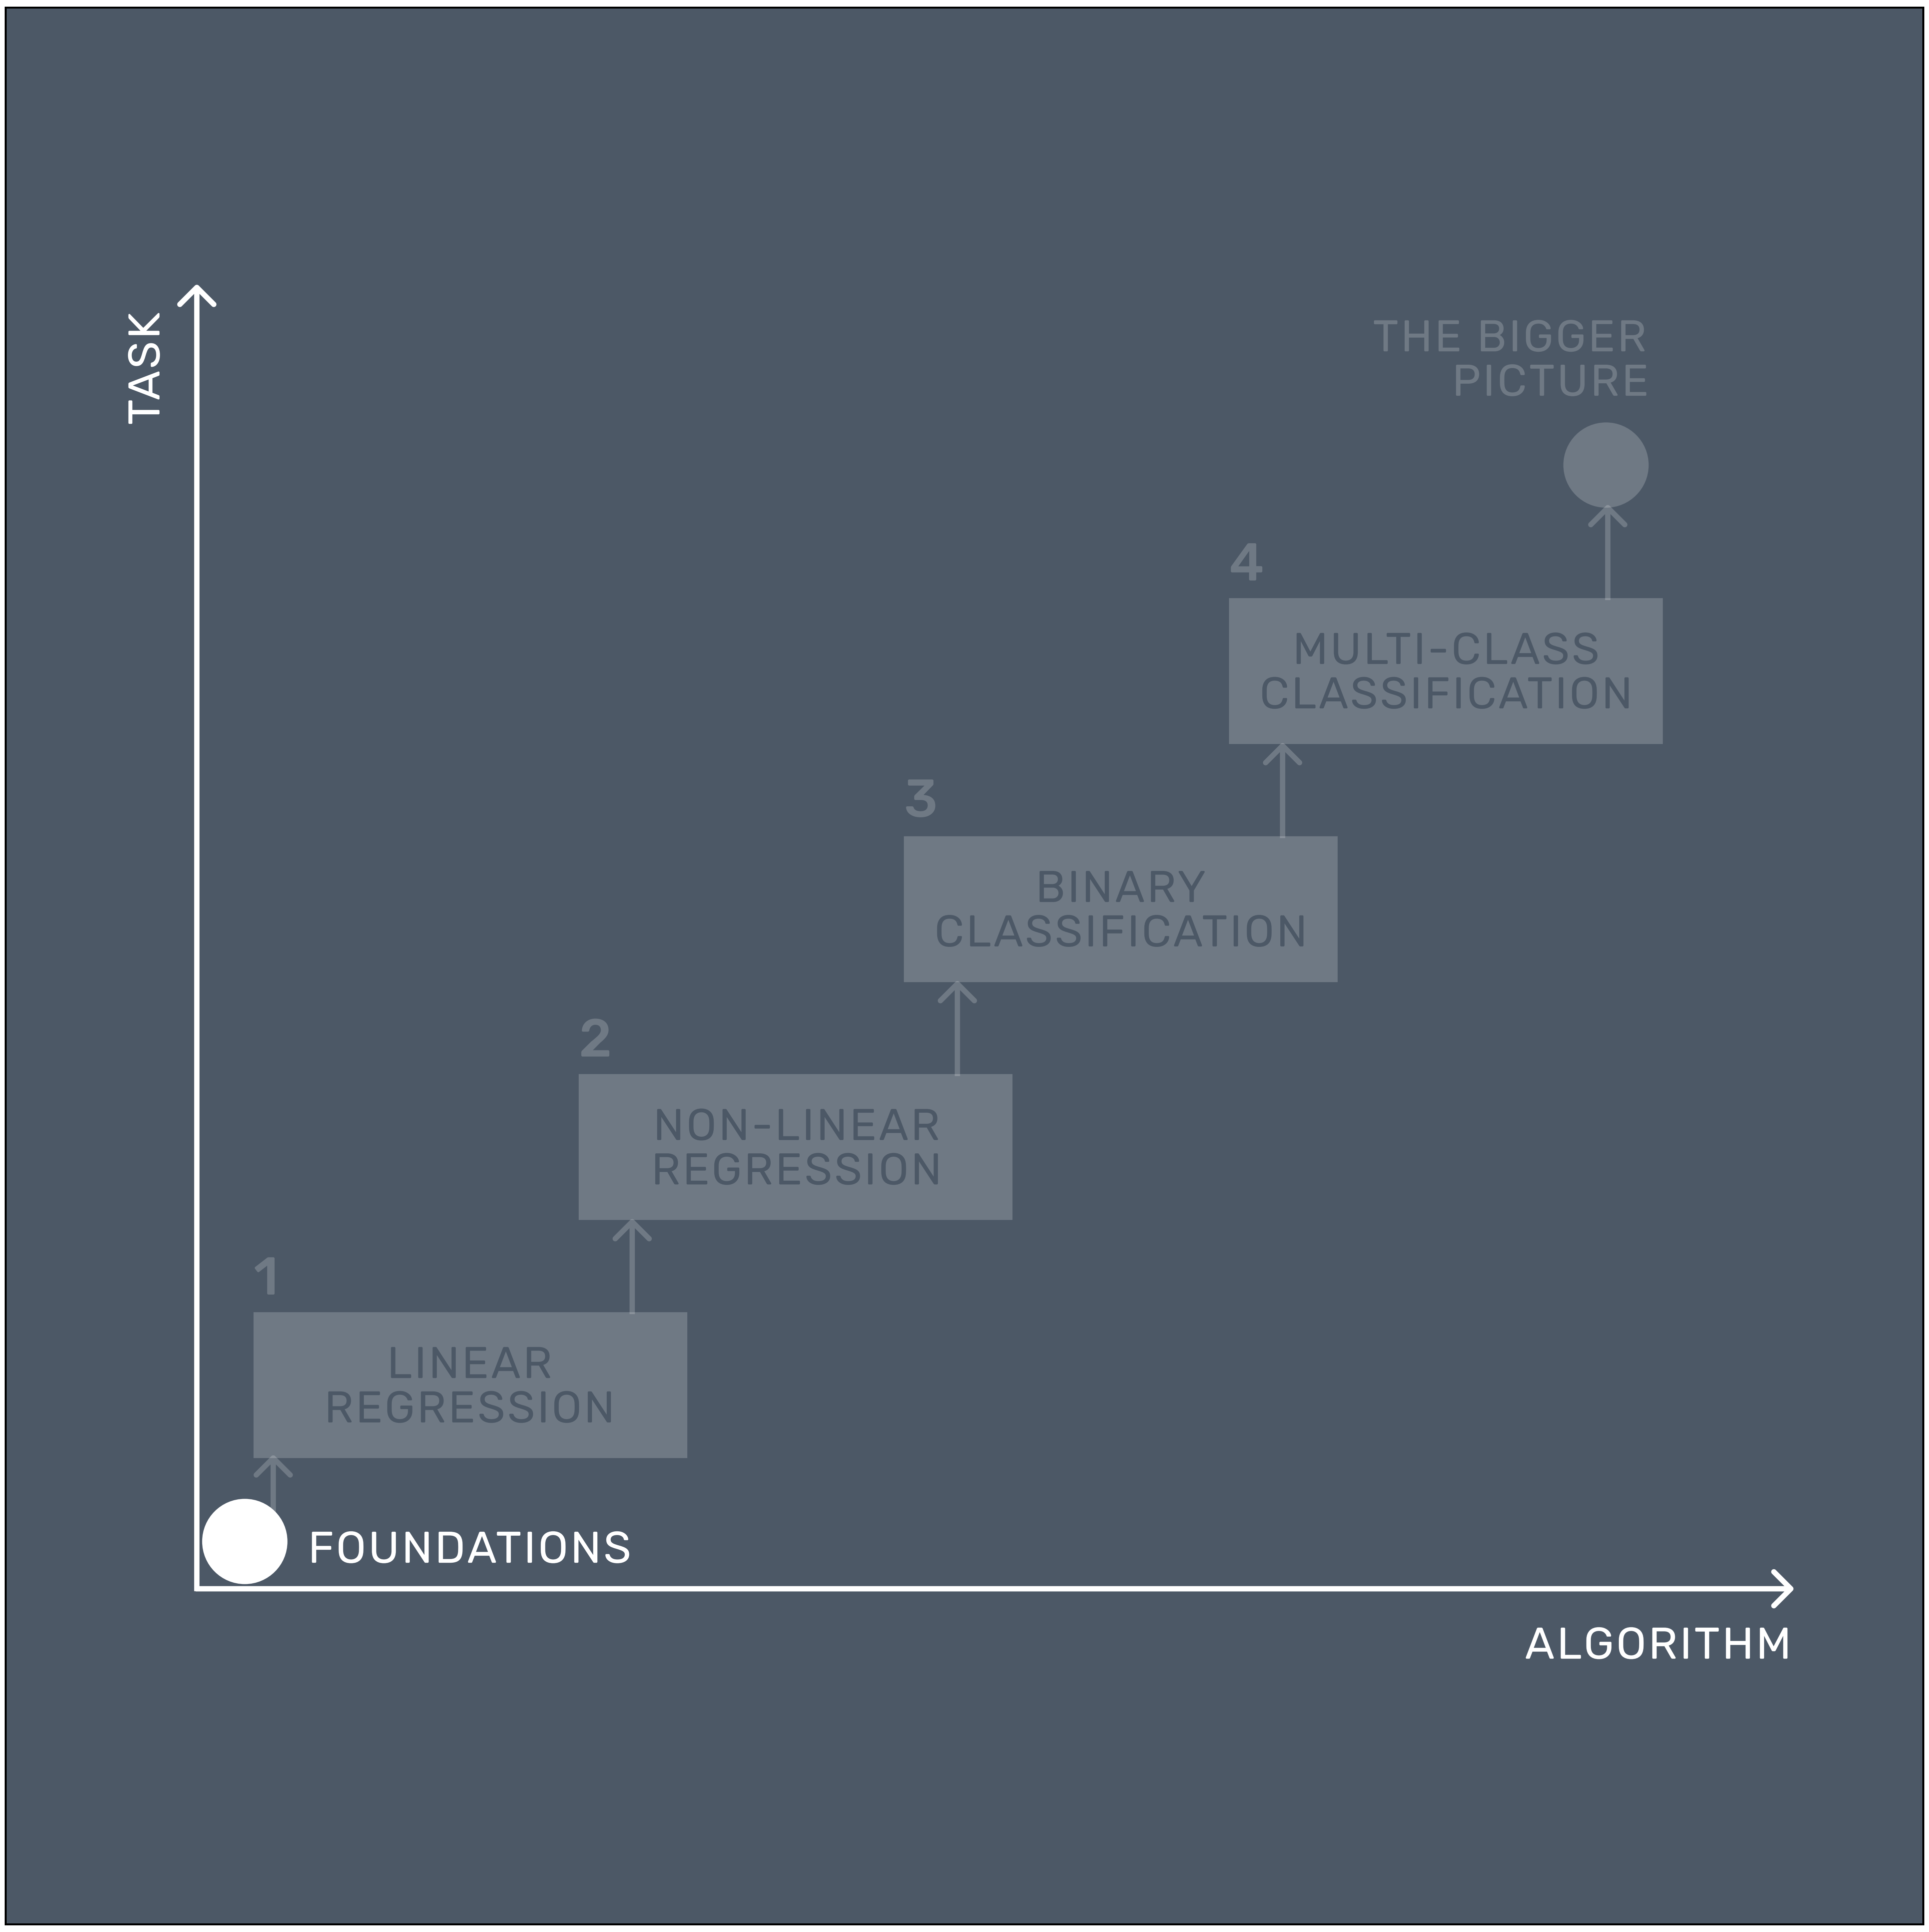

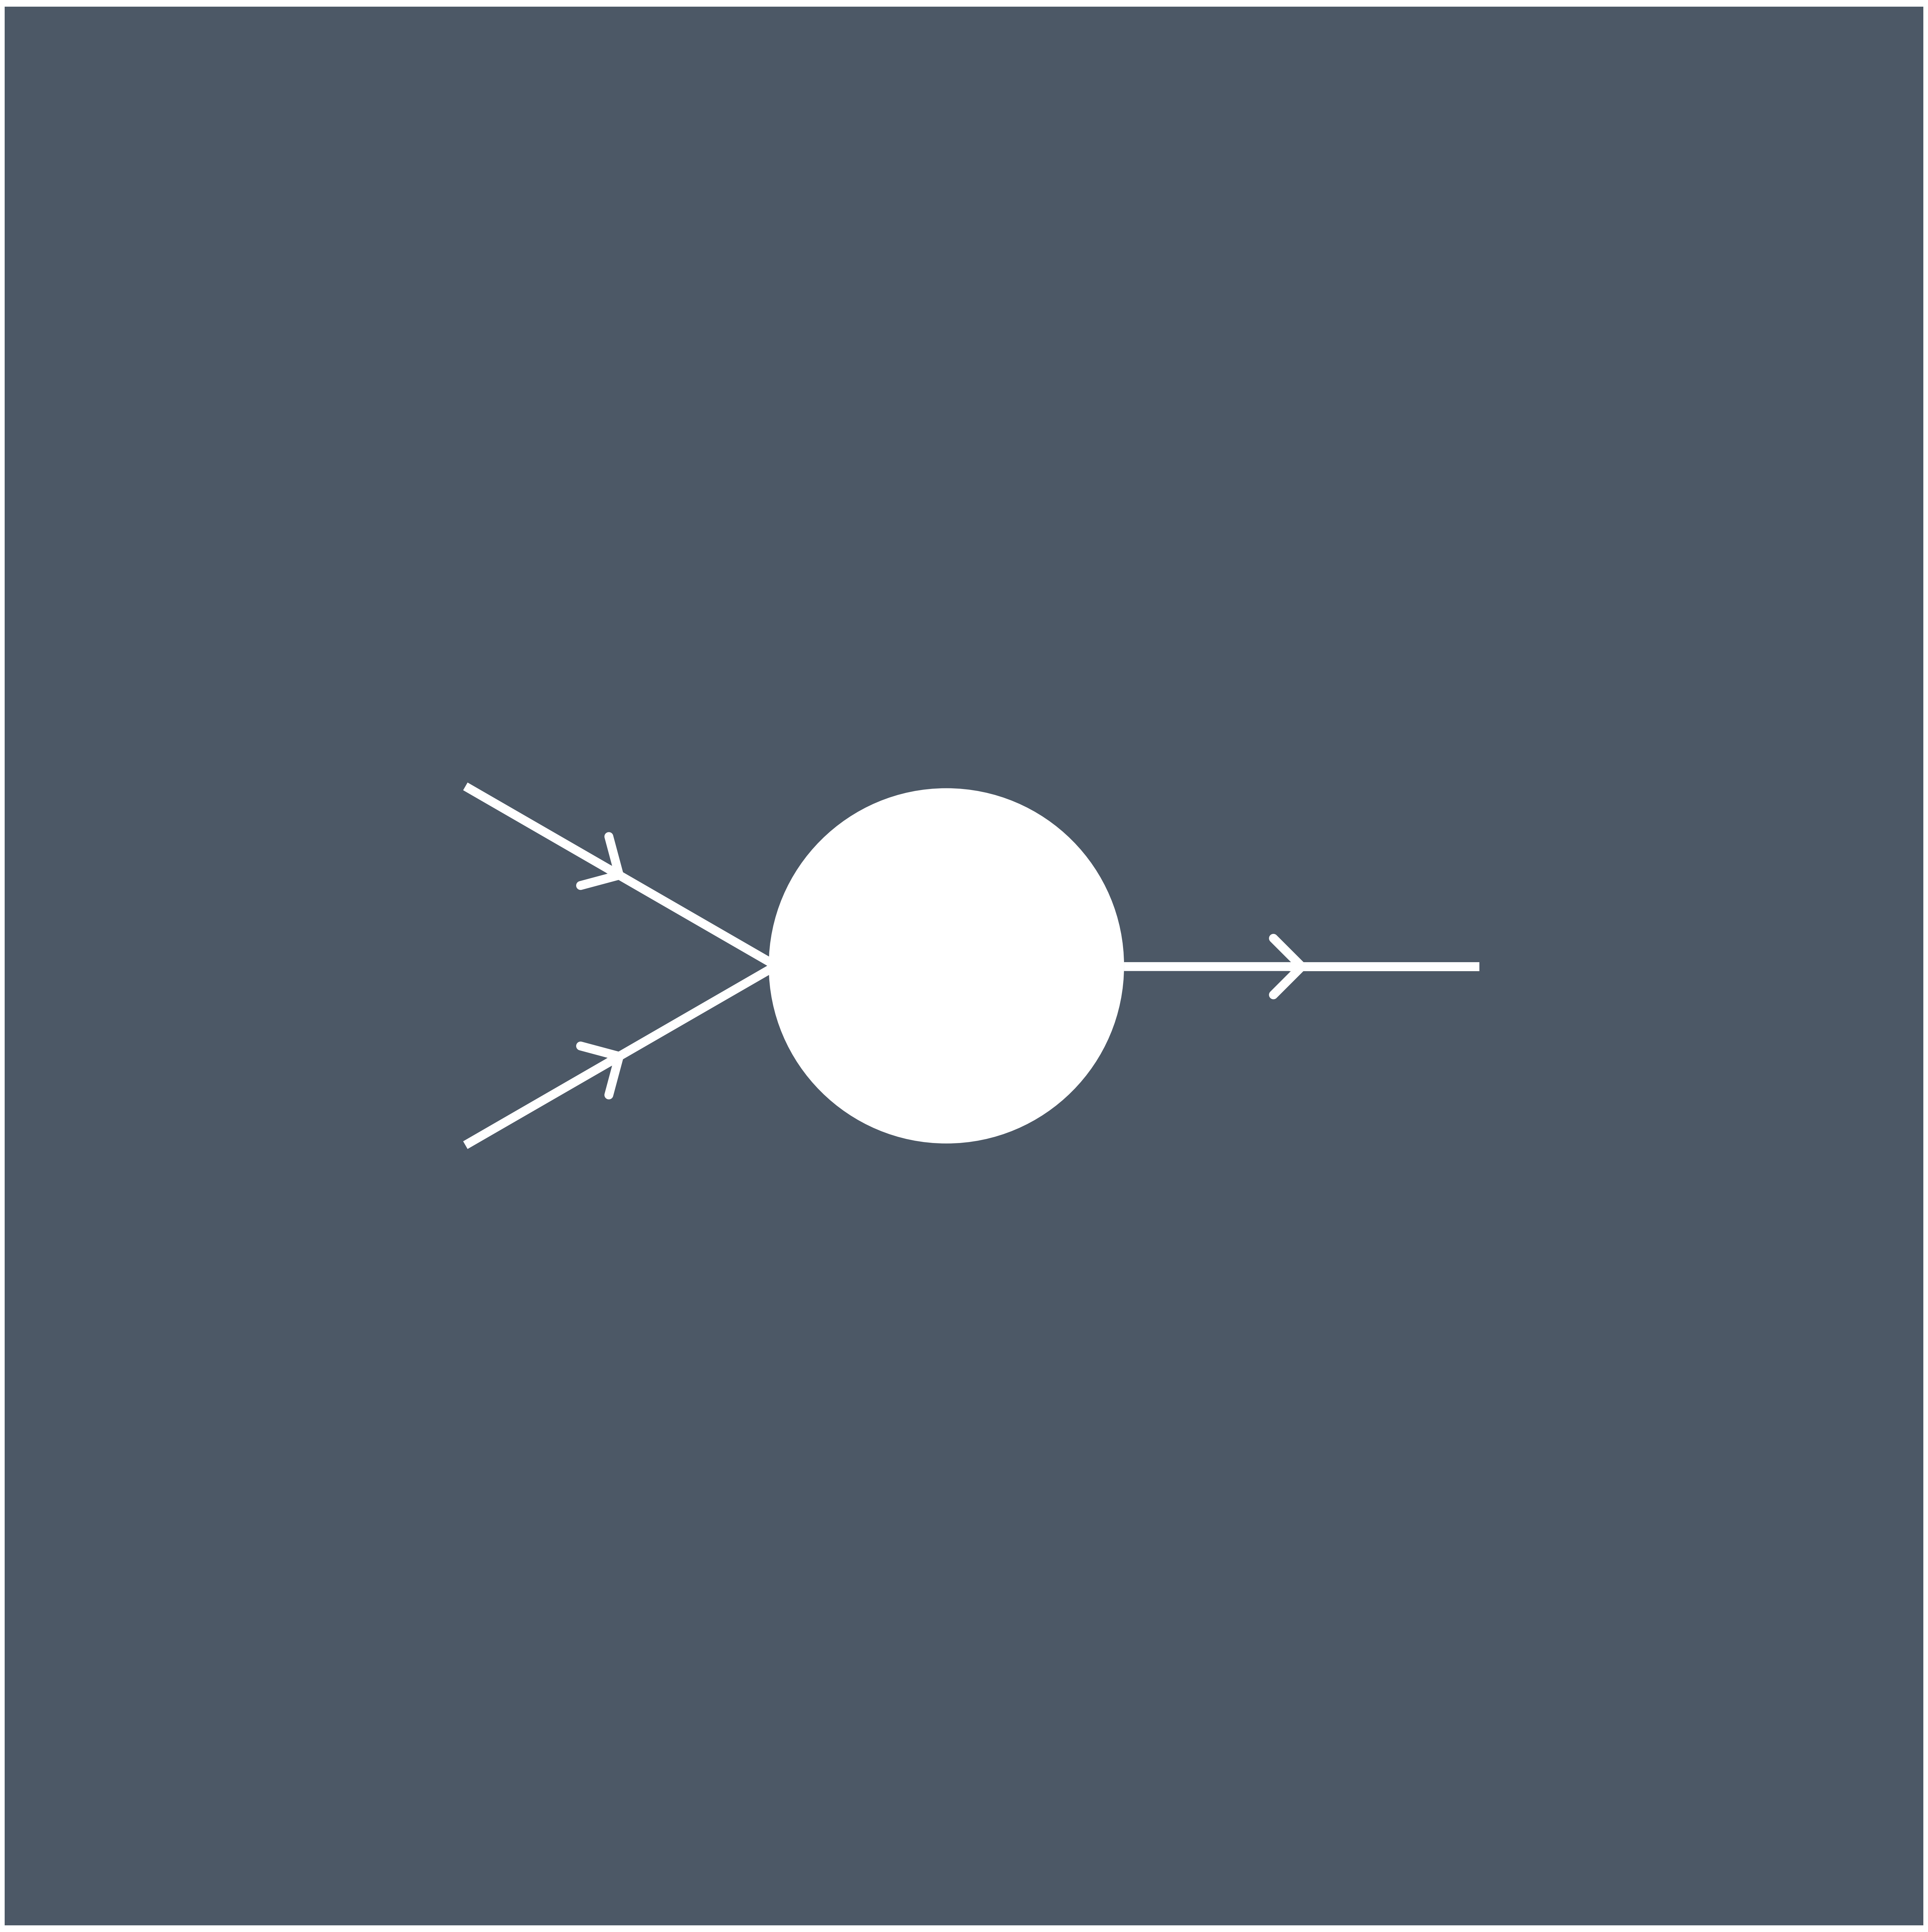

A NEURON

We have so far used the term deep learning, but from now on, we'll use *neural network* instead. These terms are used interchangeably and refer to the same thing. But as we start to go into the inner workings, neural network is a more natural term to use.

To begin our journey, let's start with a *neuron*. The neuron is the smallest unit and the building block of a neural network.

A neuron takes a set of inputs, performs some mathematical computations, and gives an output.

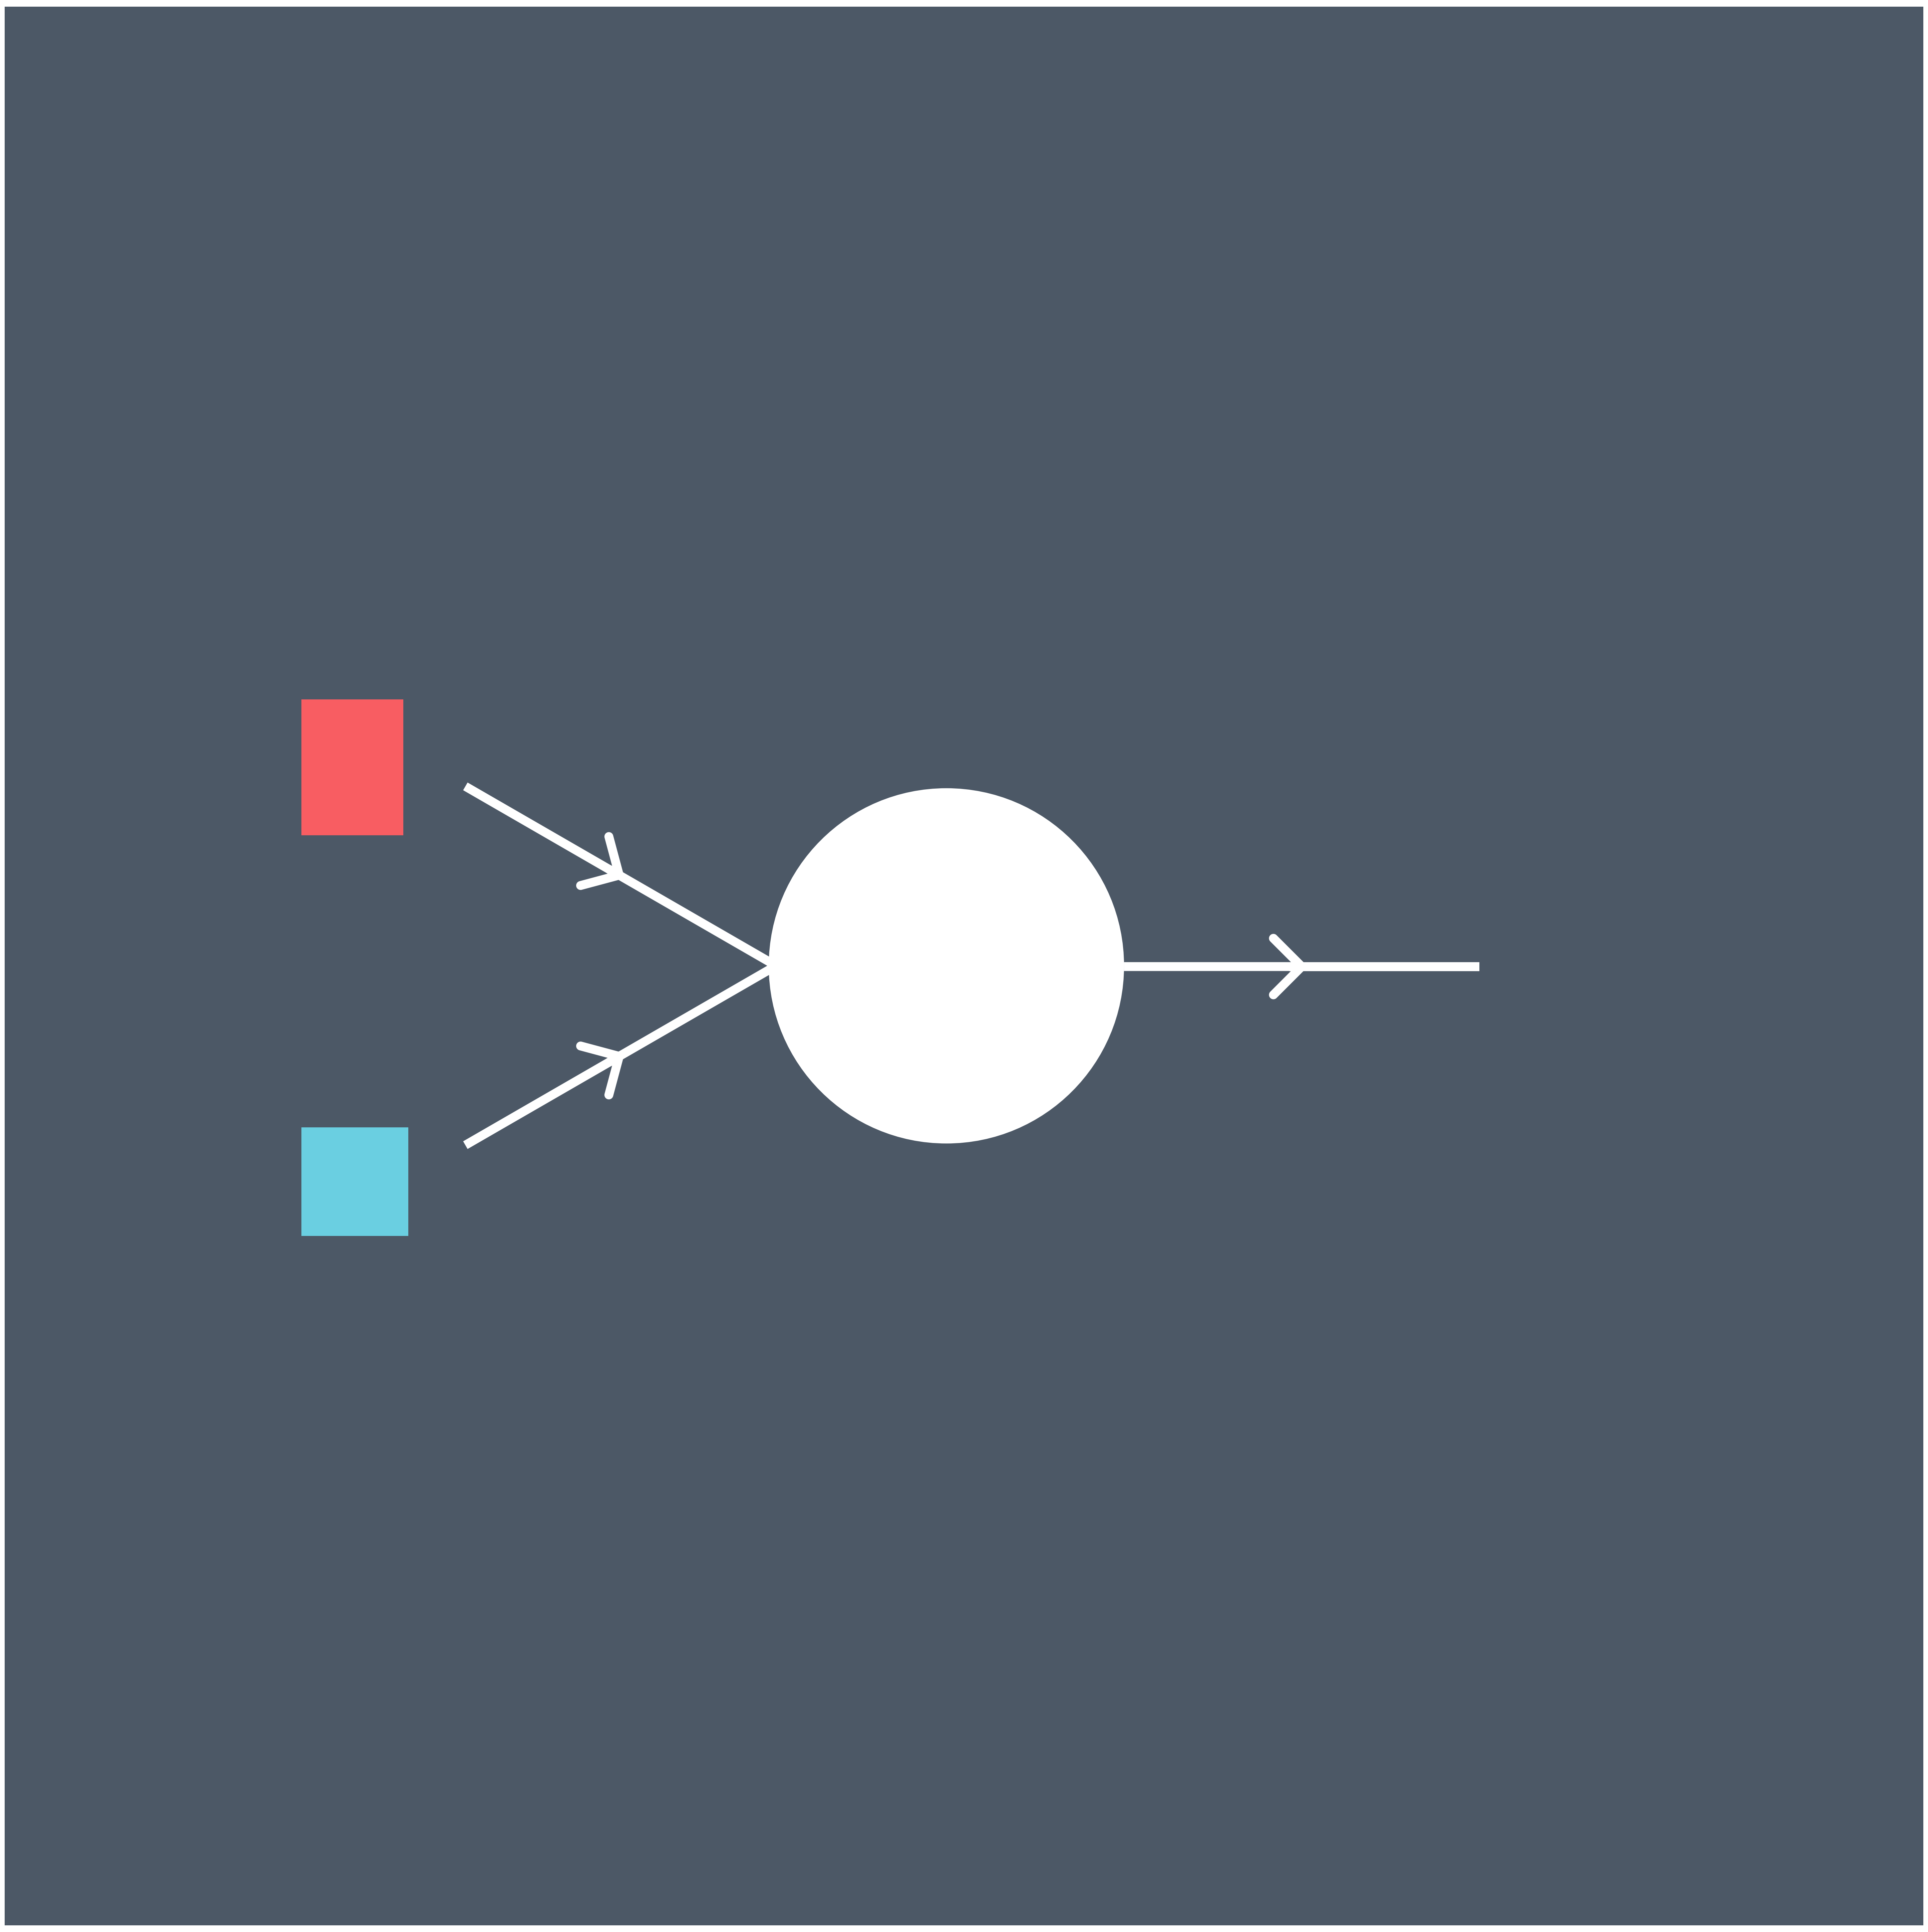

INPUTS

The inputs and outputs are numbers, either positive or negative. In this example, the neuron takes two inputs. However, there is no limit to the number of inputs a neuron can take.

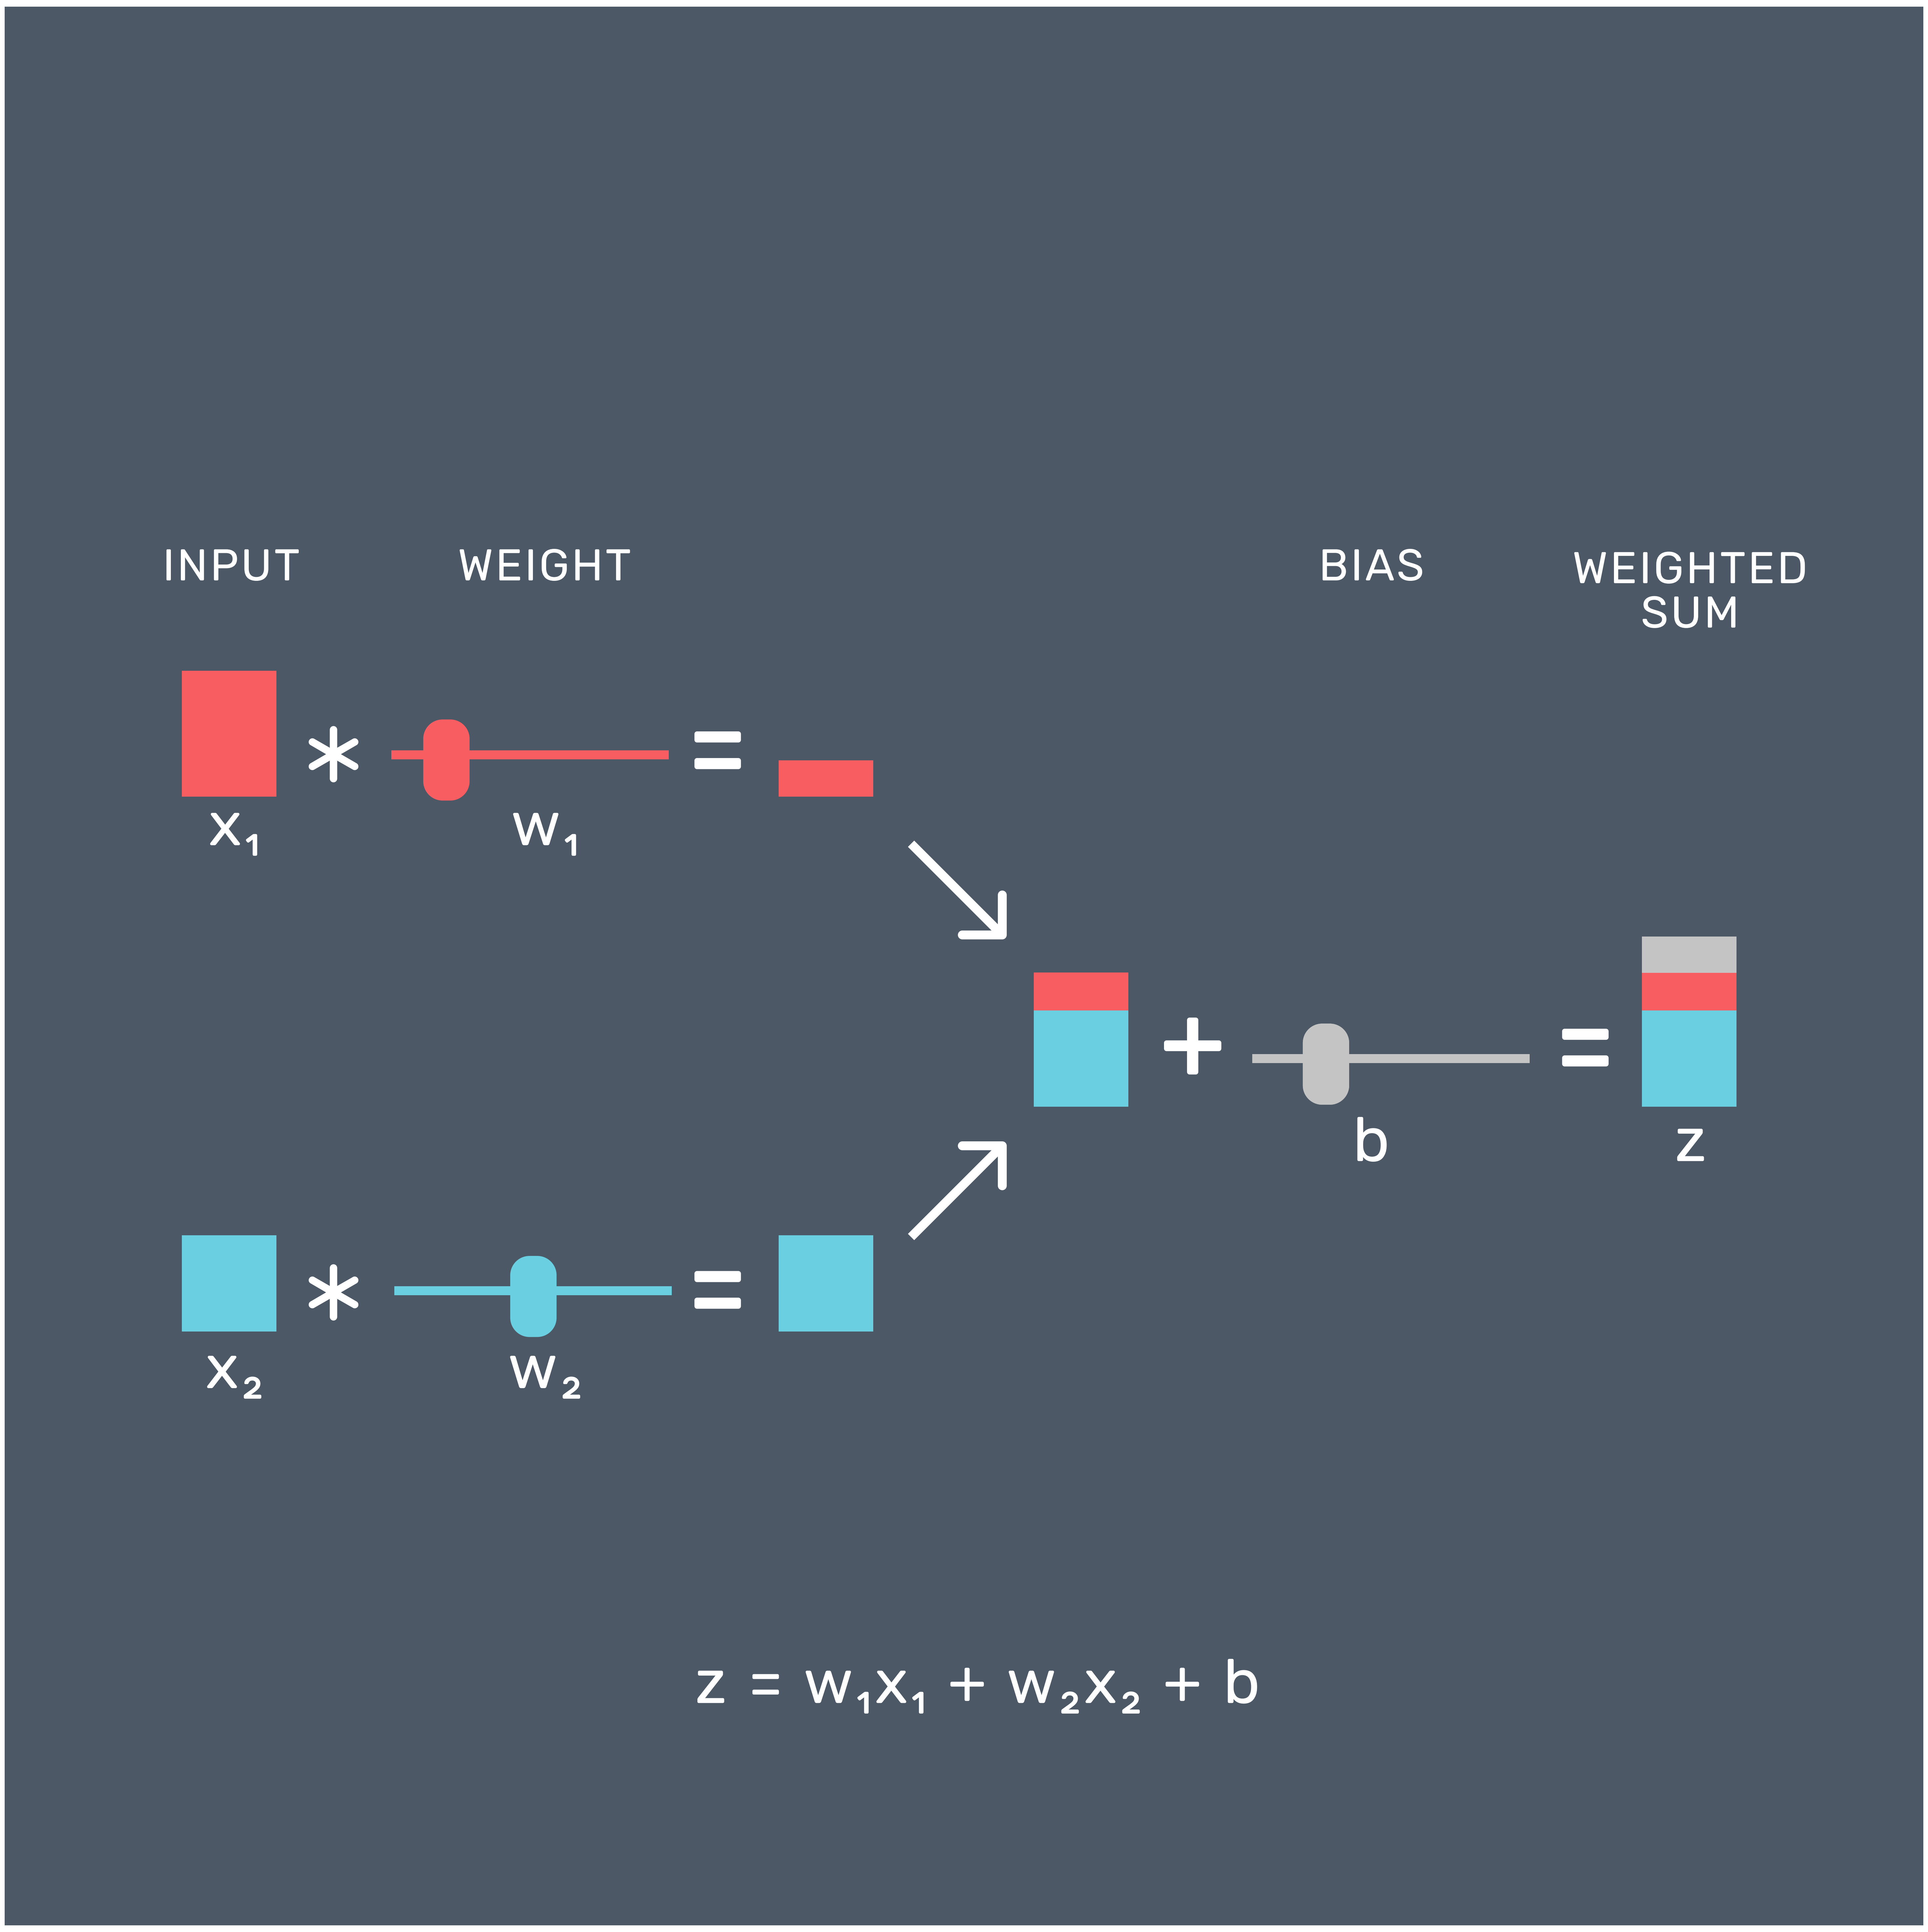

WEIGHTED SUM

The first computation that a neuron performs is the *weighted sum.* It multiplies each input by its corresponding *weight*. Then all the inputs are summed and a term called *bias* is added.

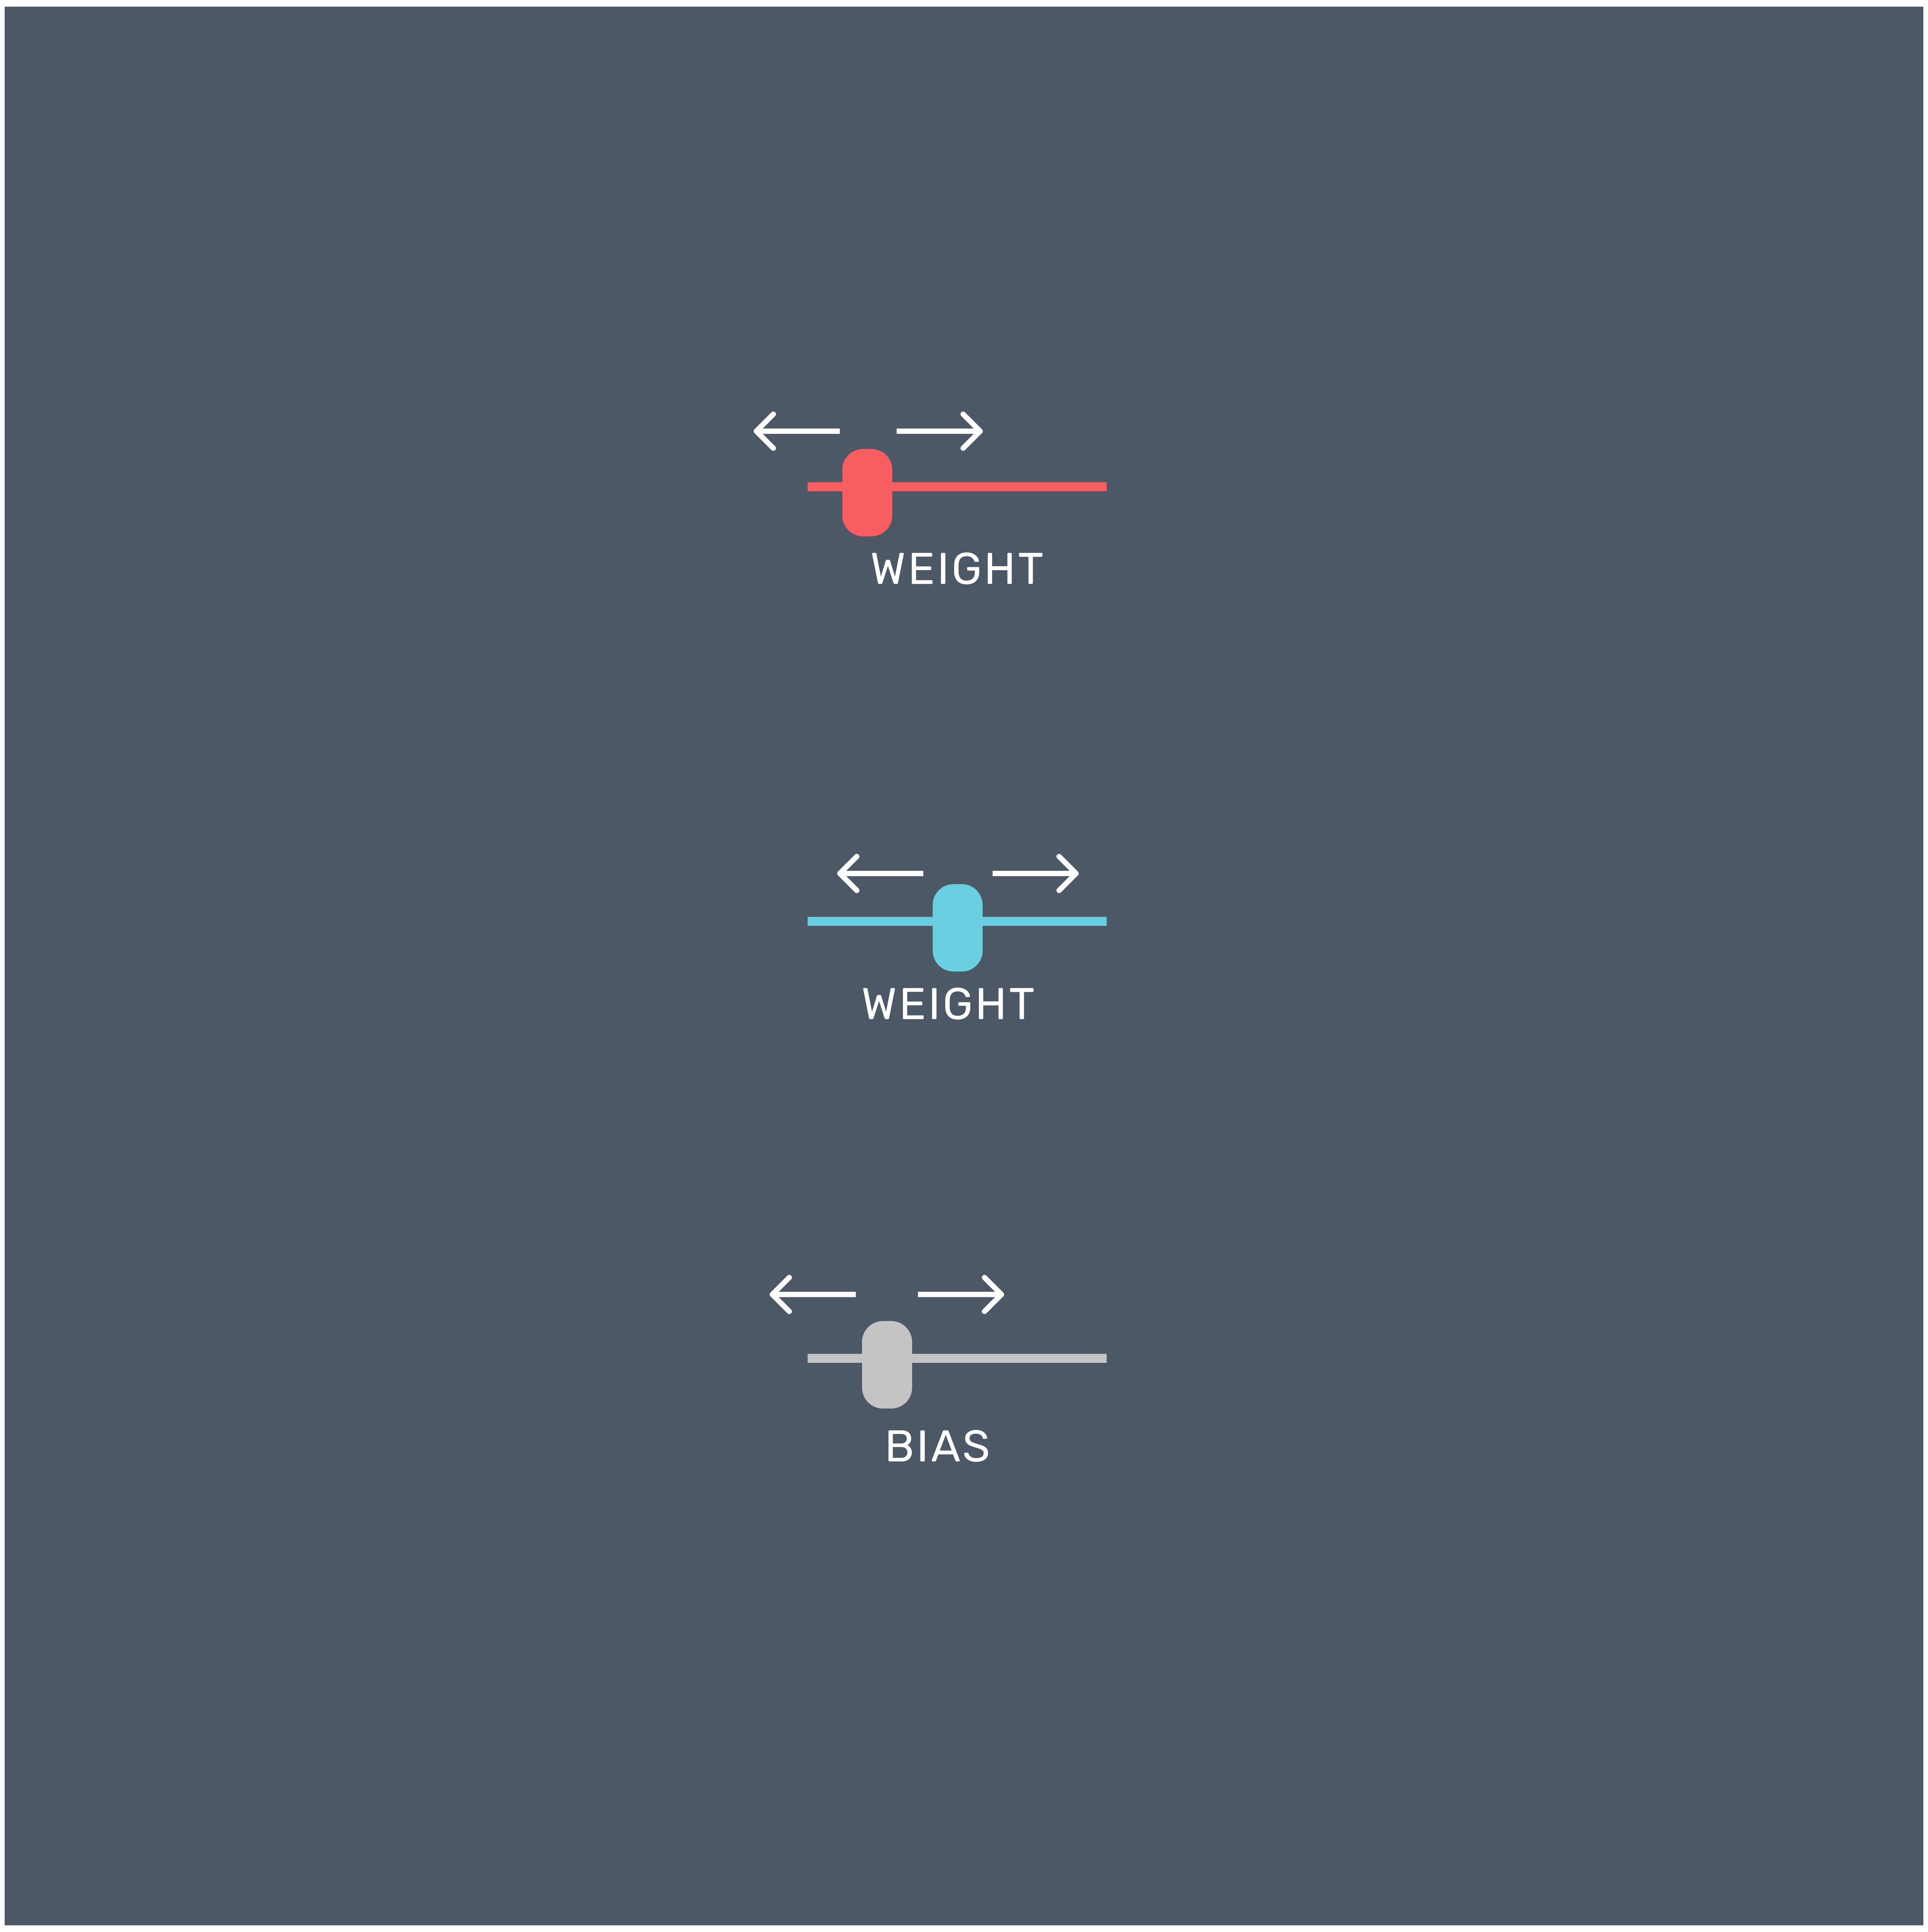

WEIGHTS AND BIASES

These weights and biases are called the *parameters* of the neural network.

These adjustable parameters are the medium through which a neural network learns, which we'll explore in detail in this book.

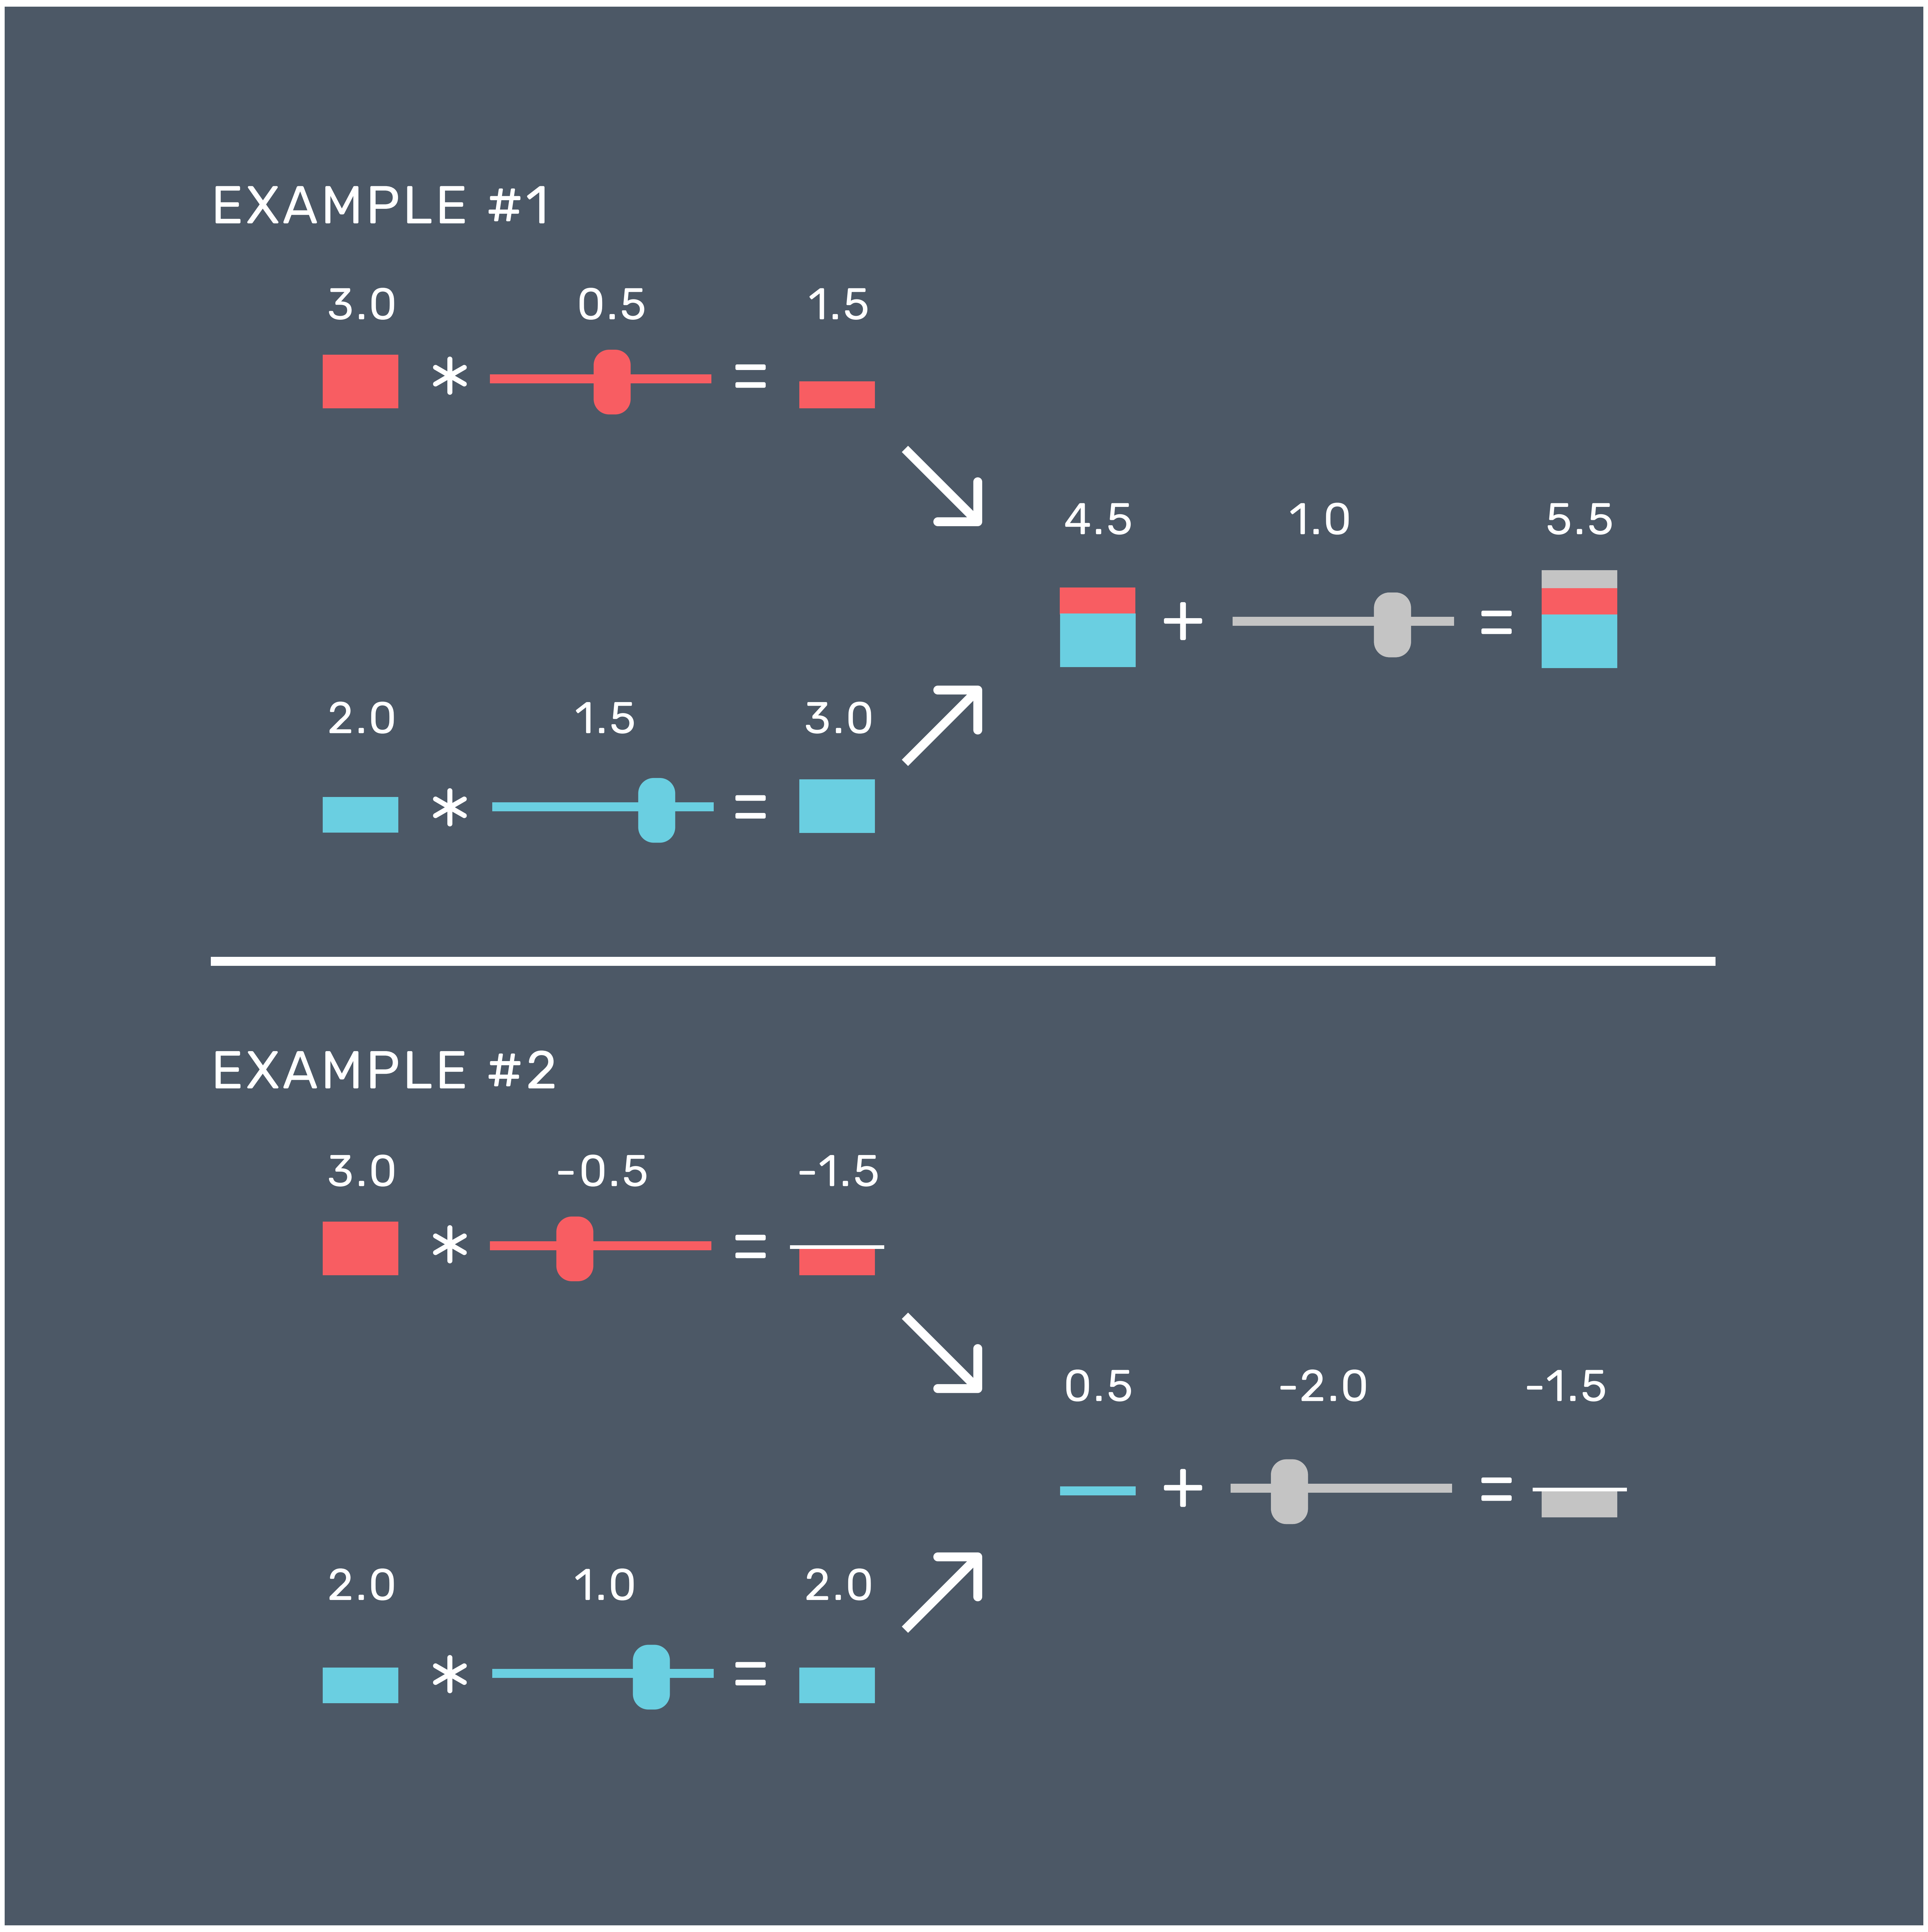

EXAMPLE

Here we have a neuron with two inputs, 3.0 and 2.0. Given different weight values, it will correspondingly output different values.

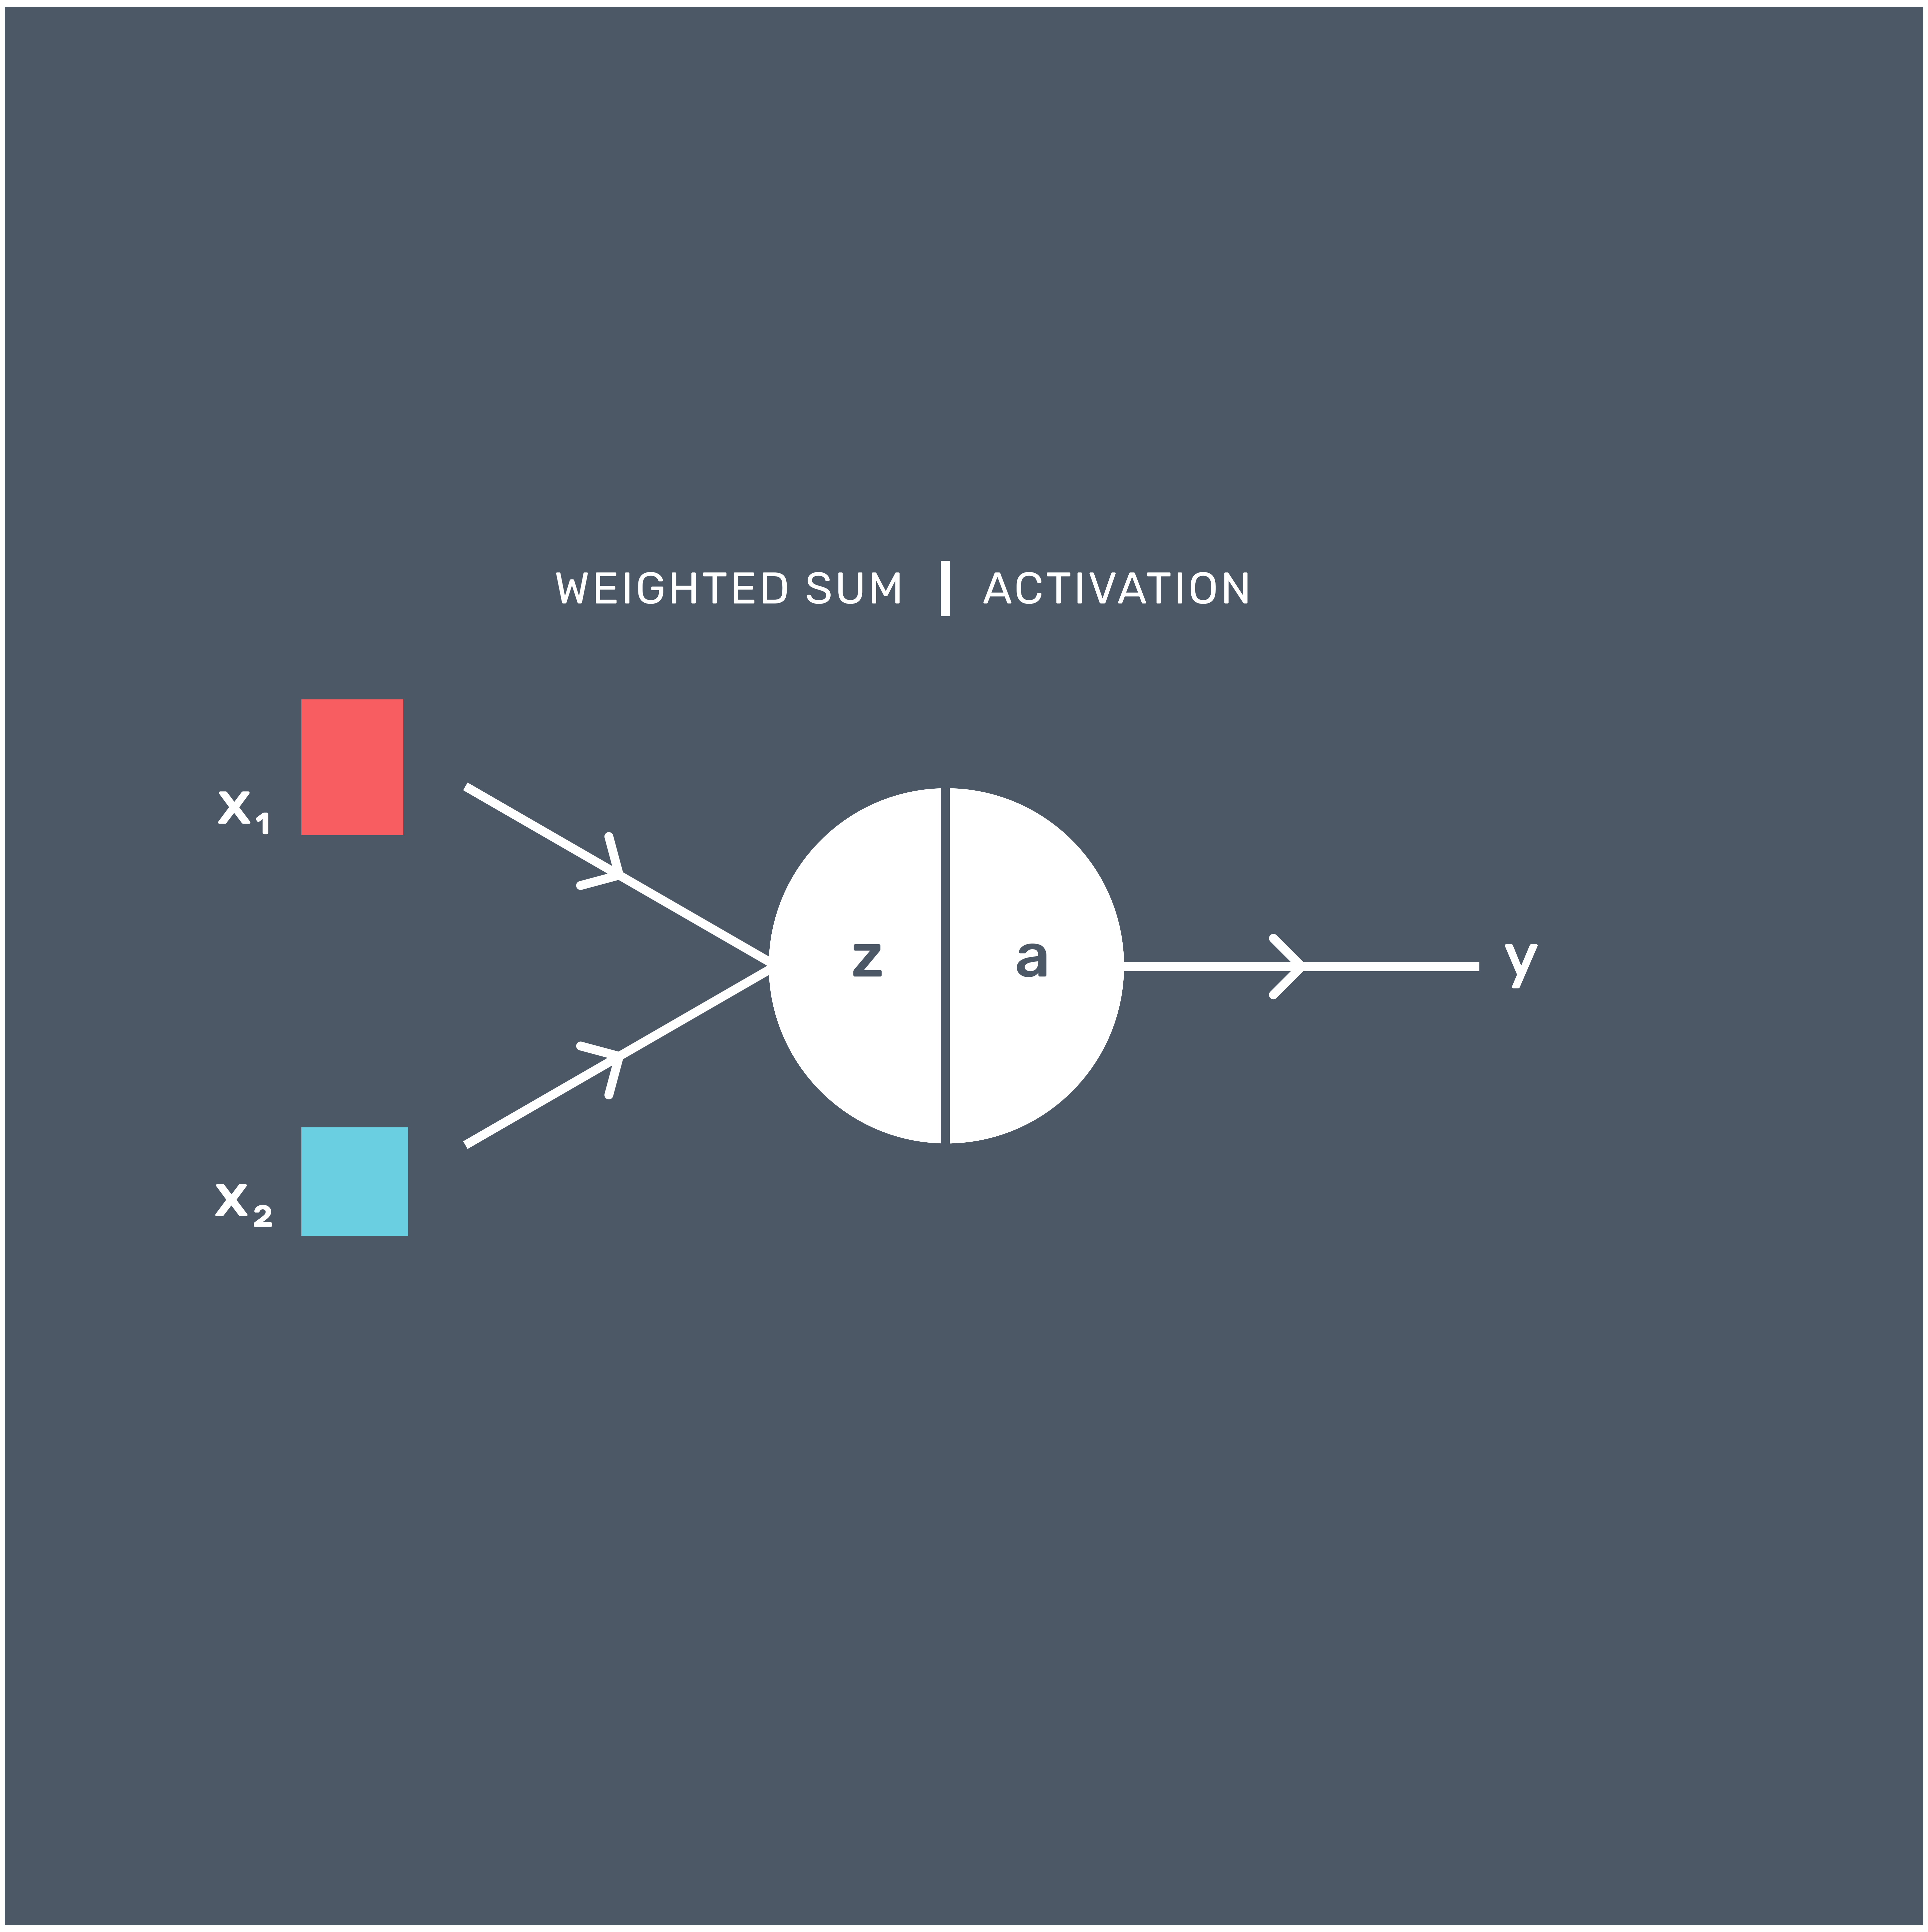

ACTIVATION

The second computation performed by the neuron is called an *activation*. This is done by taking the output of the weighted sum and passing it through an *activation function*.

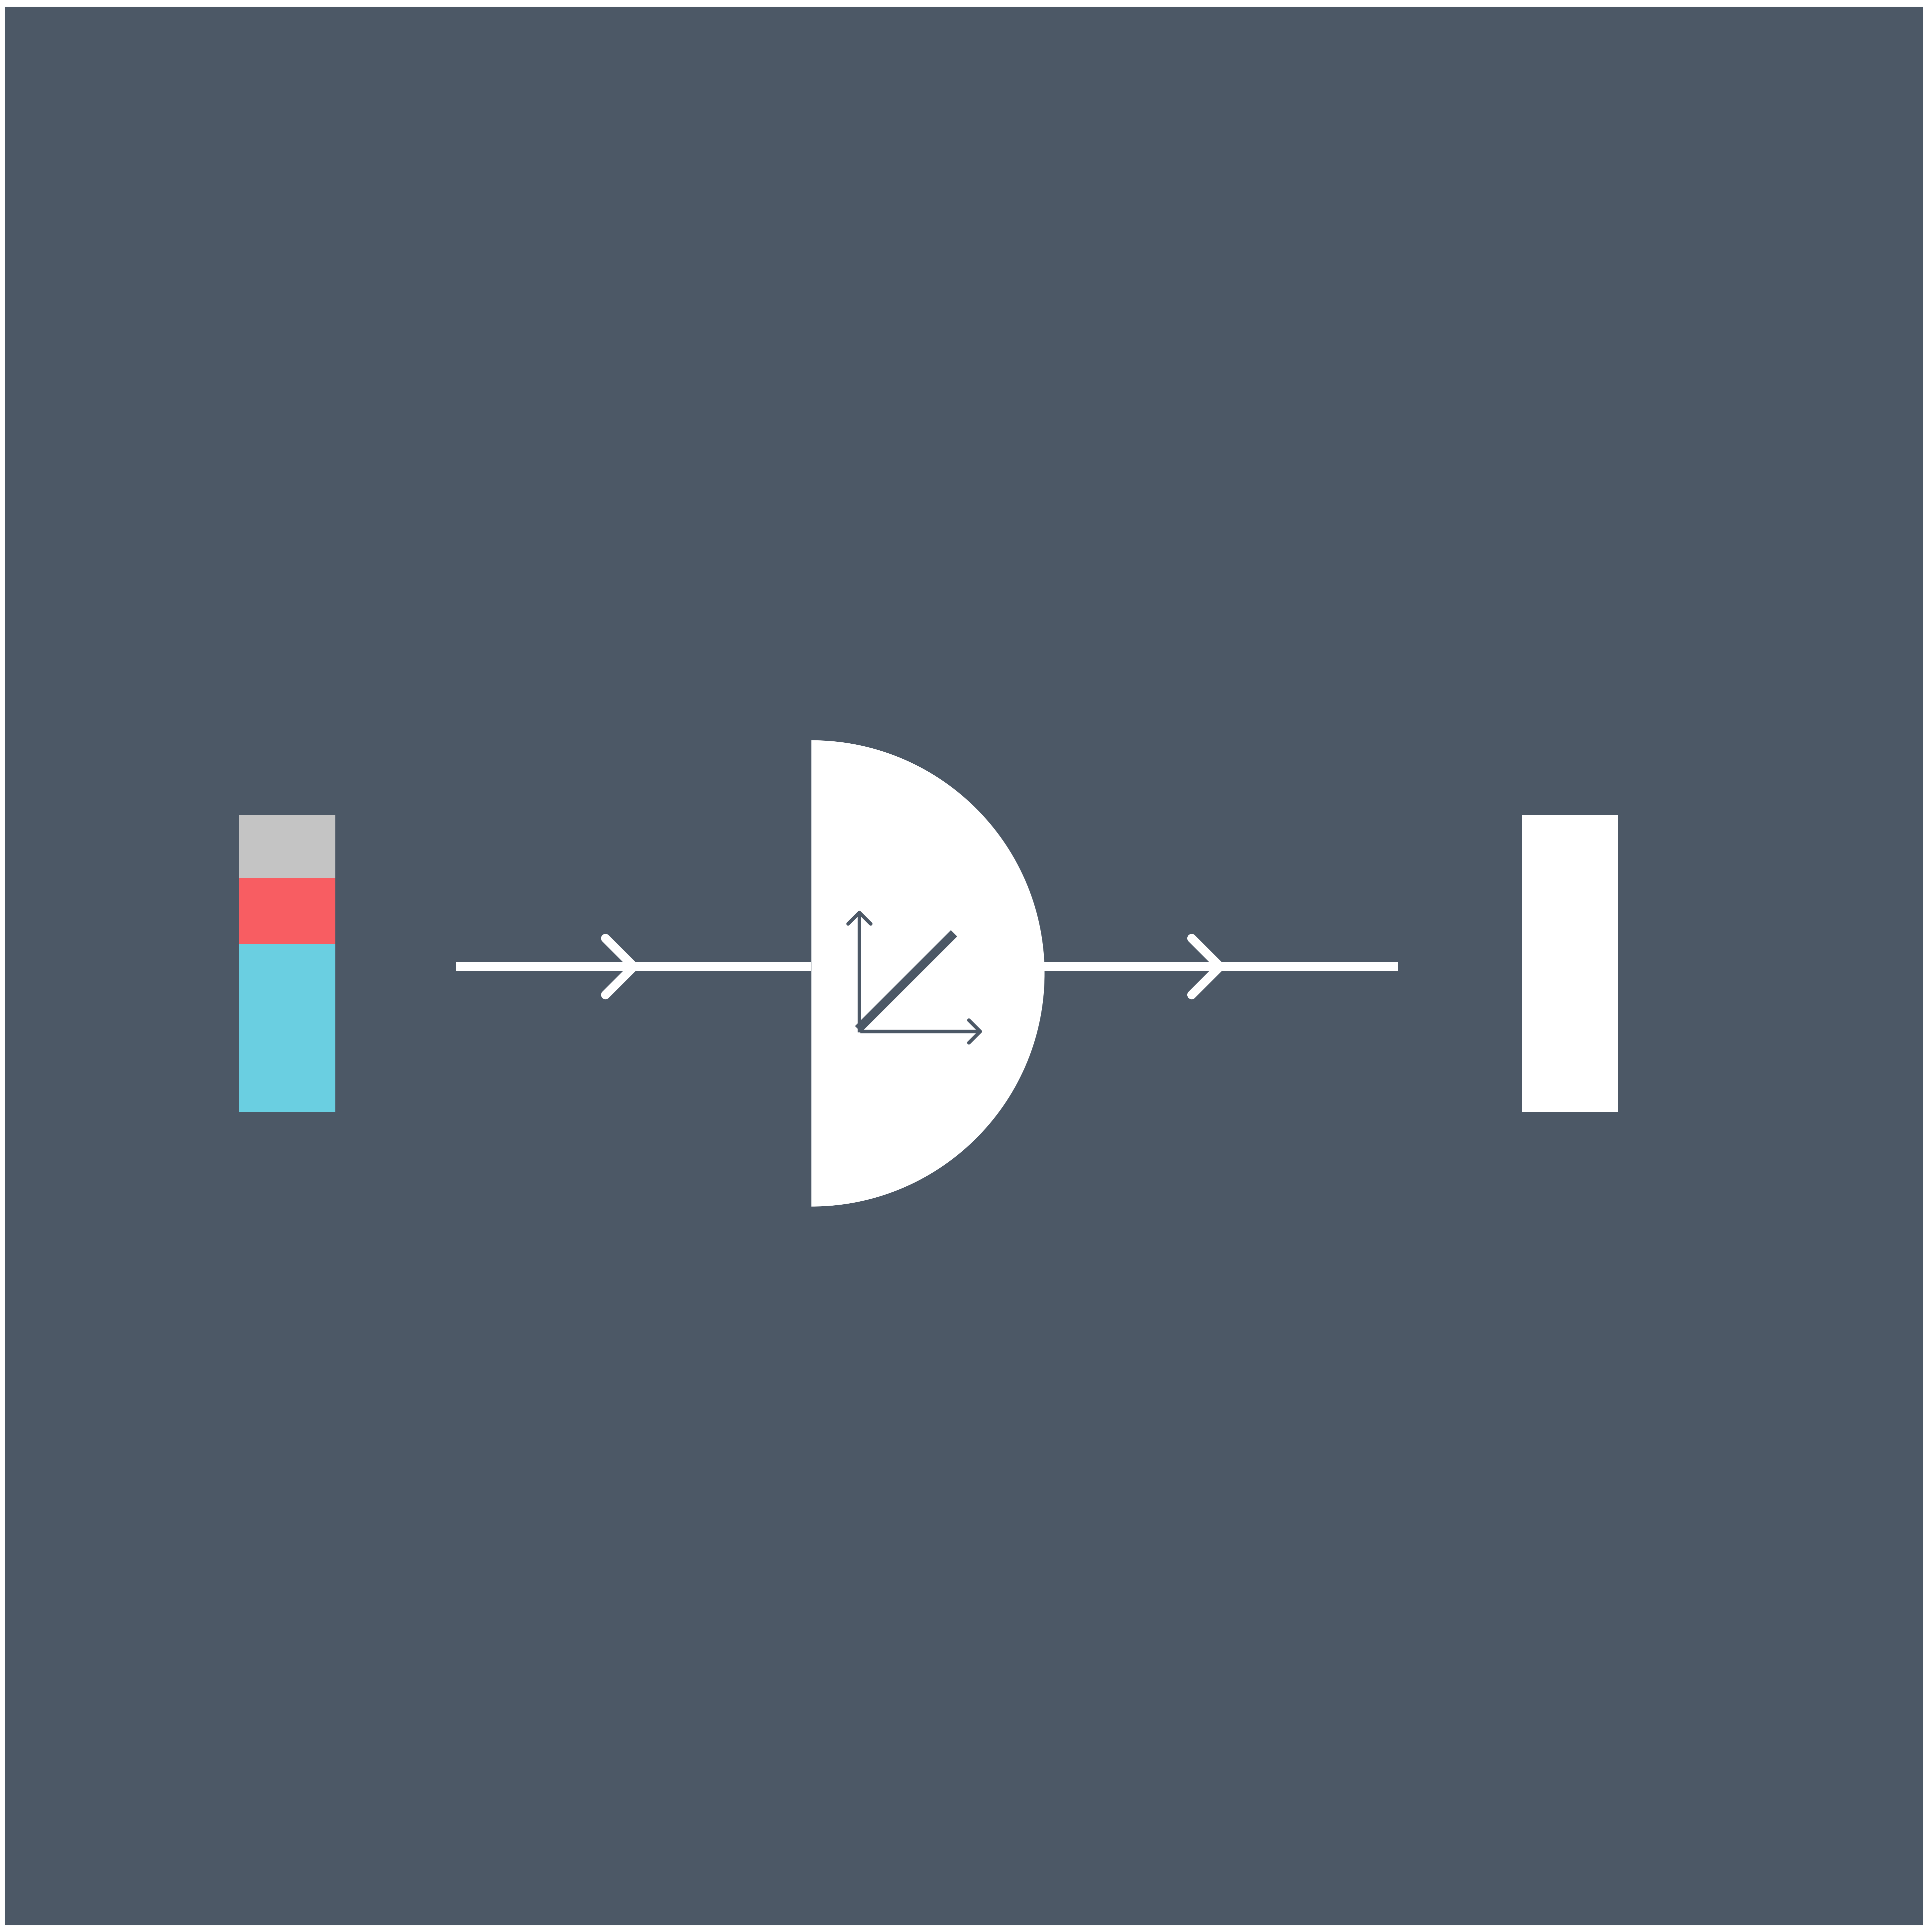

LINEAR ACTIVATION

The activation function gives the neural network the ability to express itself. This will not make much sense now but will become clear by the end of this book.

There are a few common activation functions. To start with, here we have a linear activation function. It's as basic as it gets - it simply outputs the same input it receives. Plotted on a chart, it gives a straight-line relationship between the input and the output.

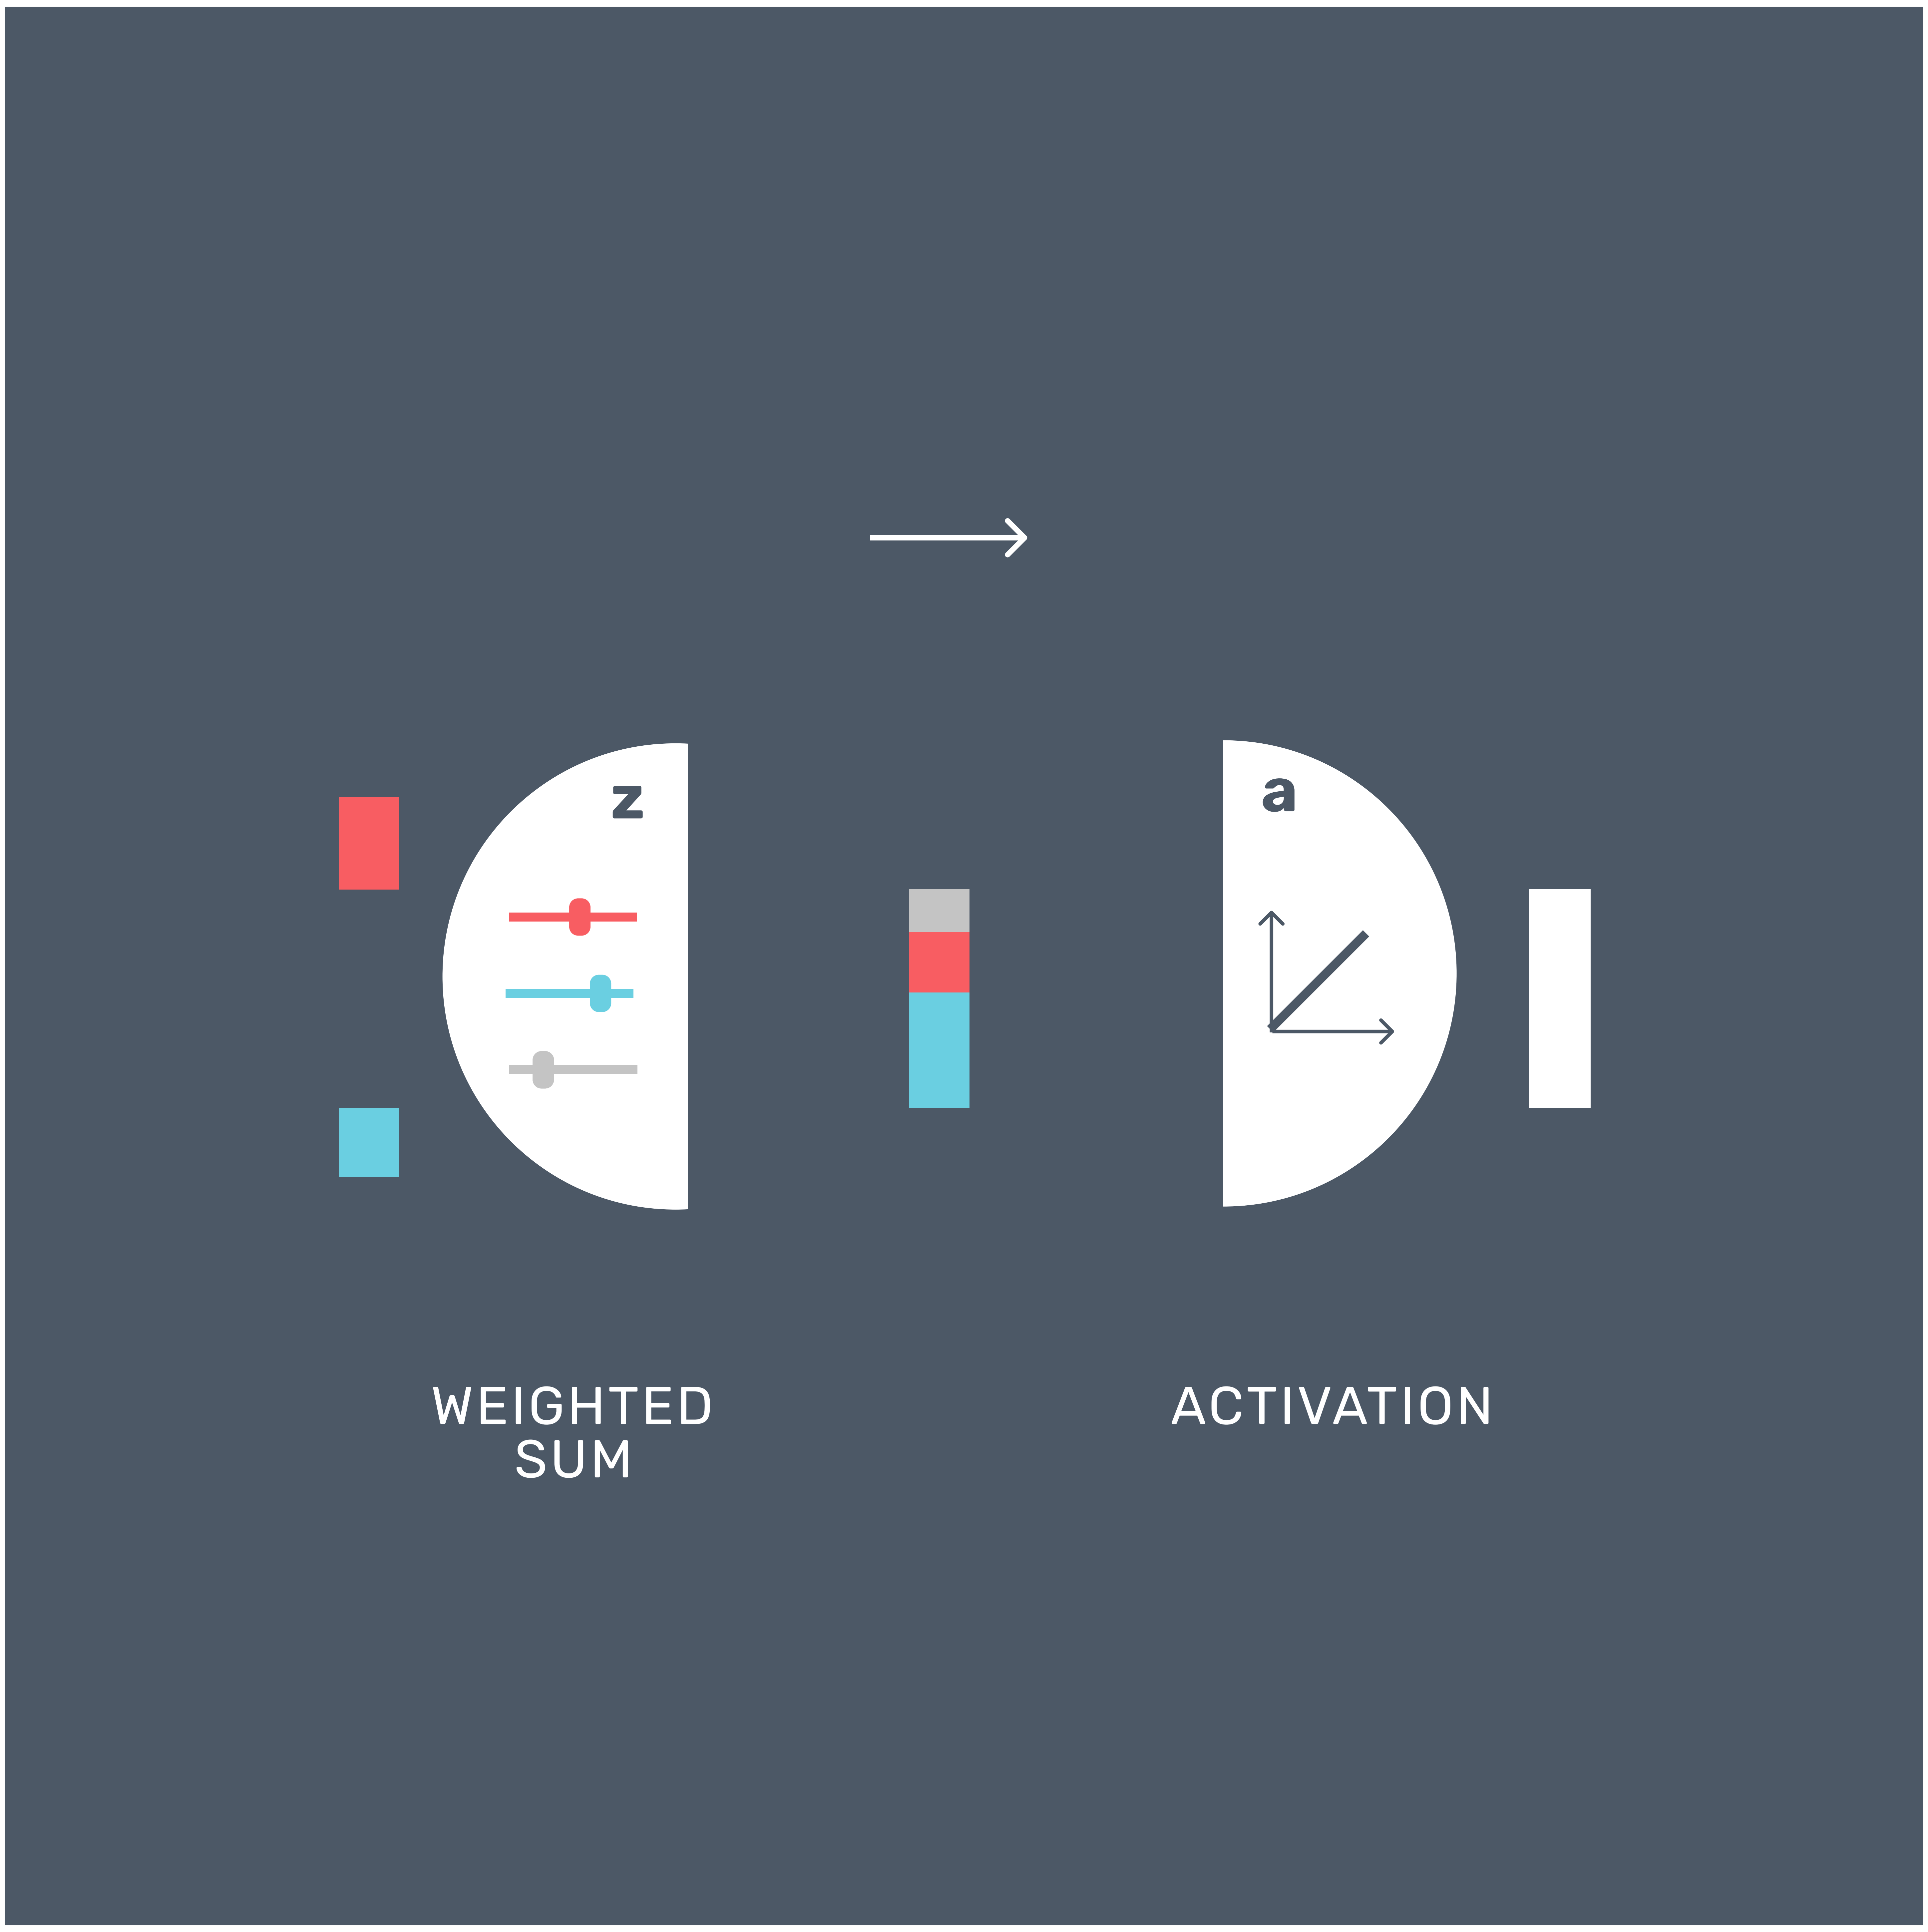

RECAP

Let's do a quick recap. When inputs are passed through a neuron, it performs a sequence of computations.

First it performs a weighted sum by multiplying each input by its corresponding weight, summing them, and finally adding a bias term.

Then it performs an activation via an activation function, which in our case, is a linear function.

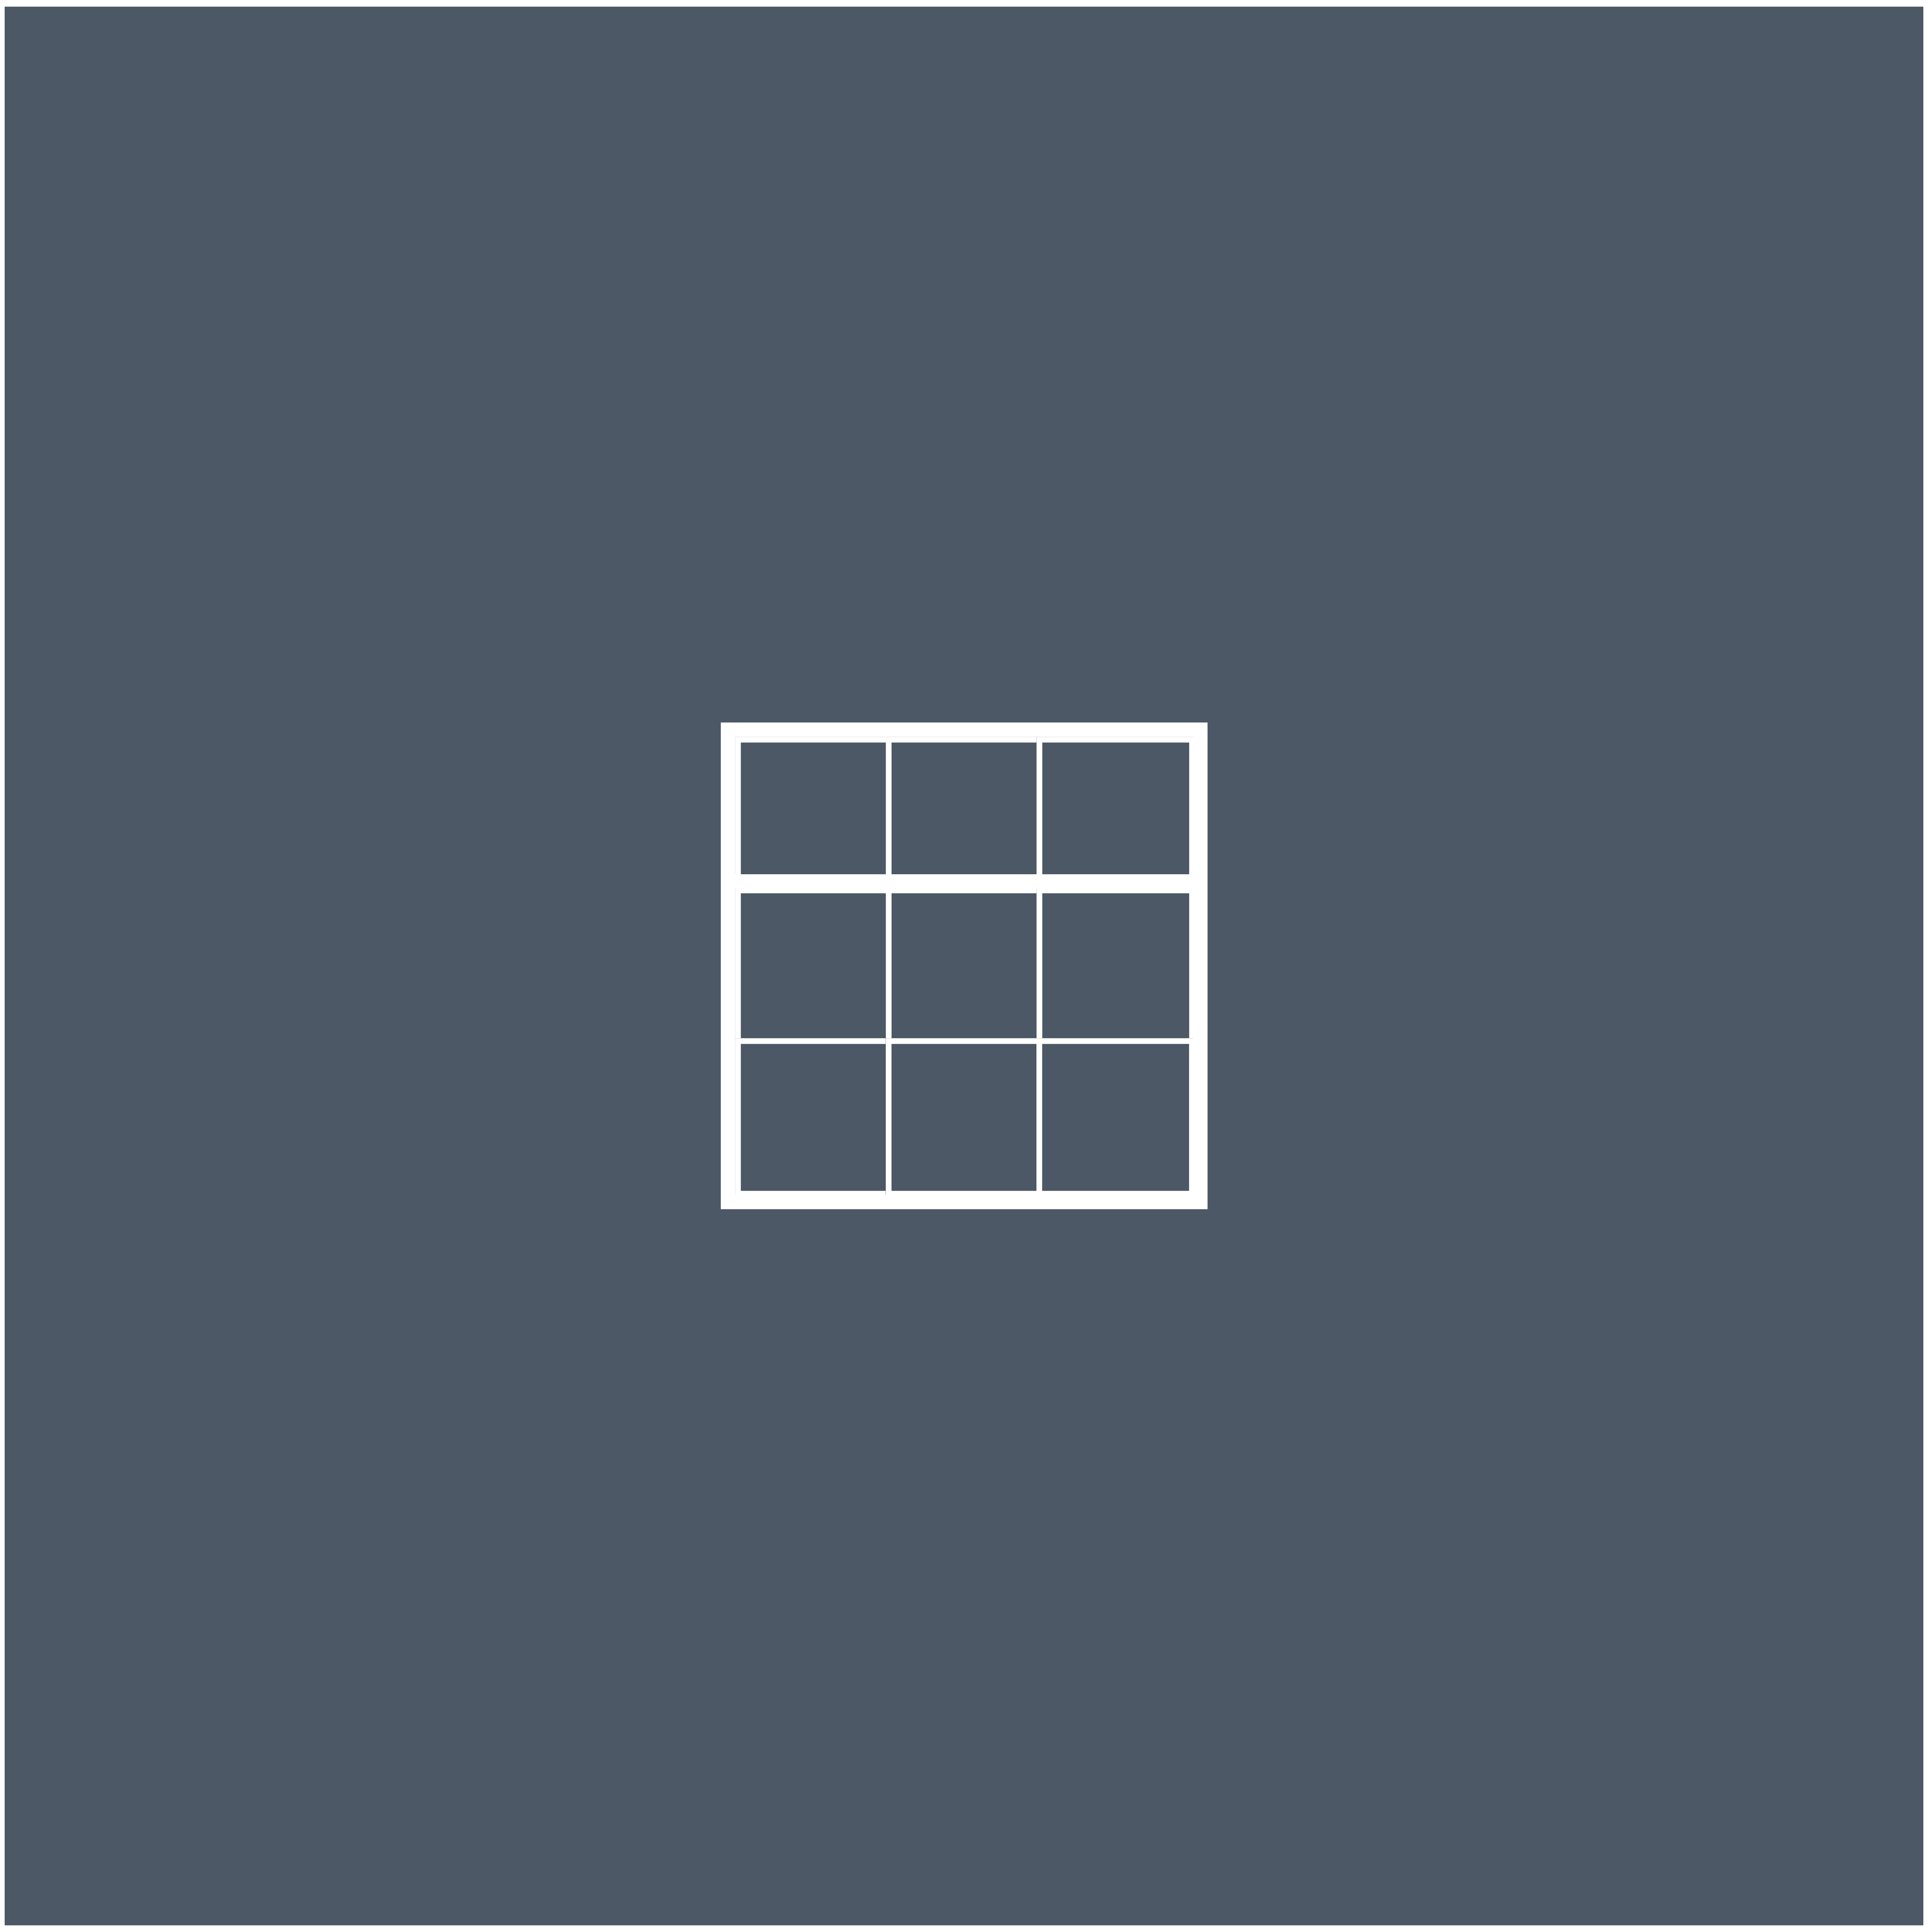

DATA

Neural networks are nothing without data. Let’s now turn our attention to data and what it brings.

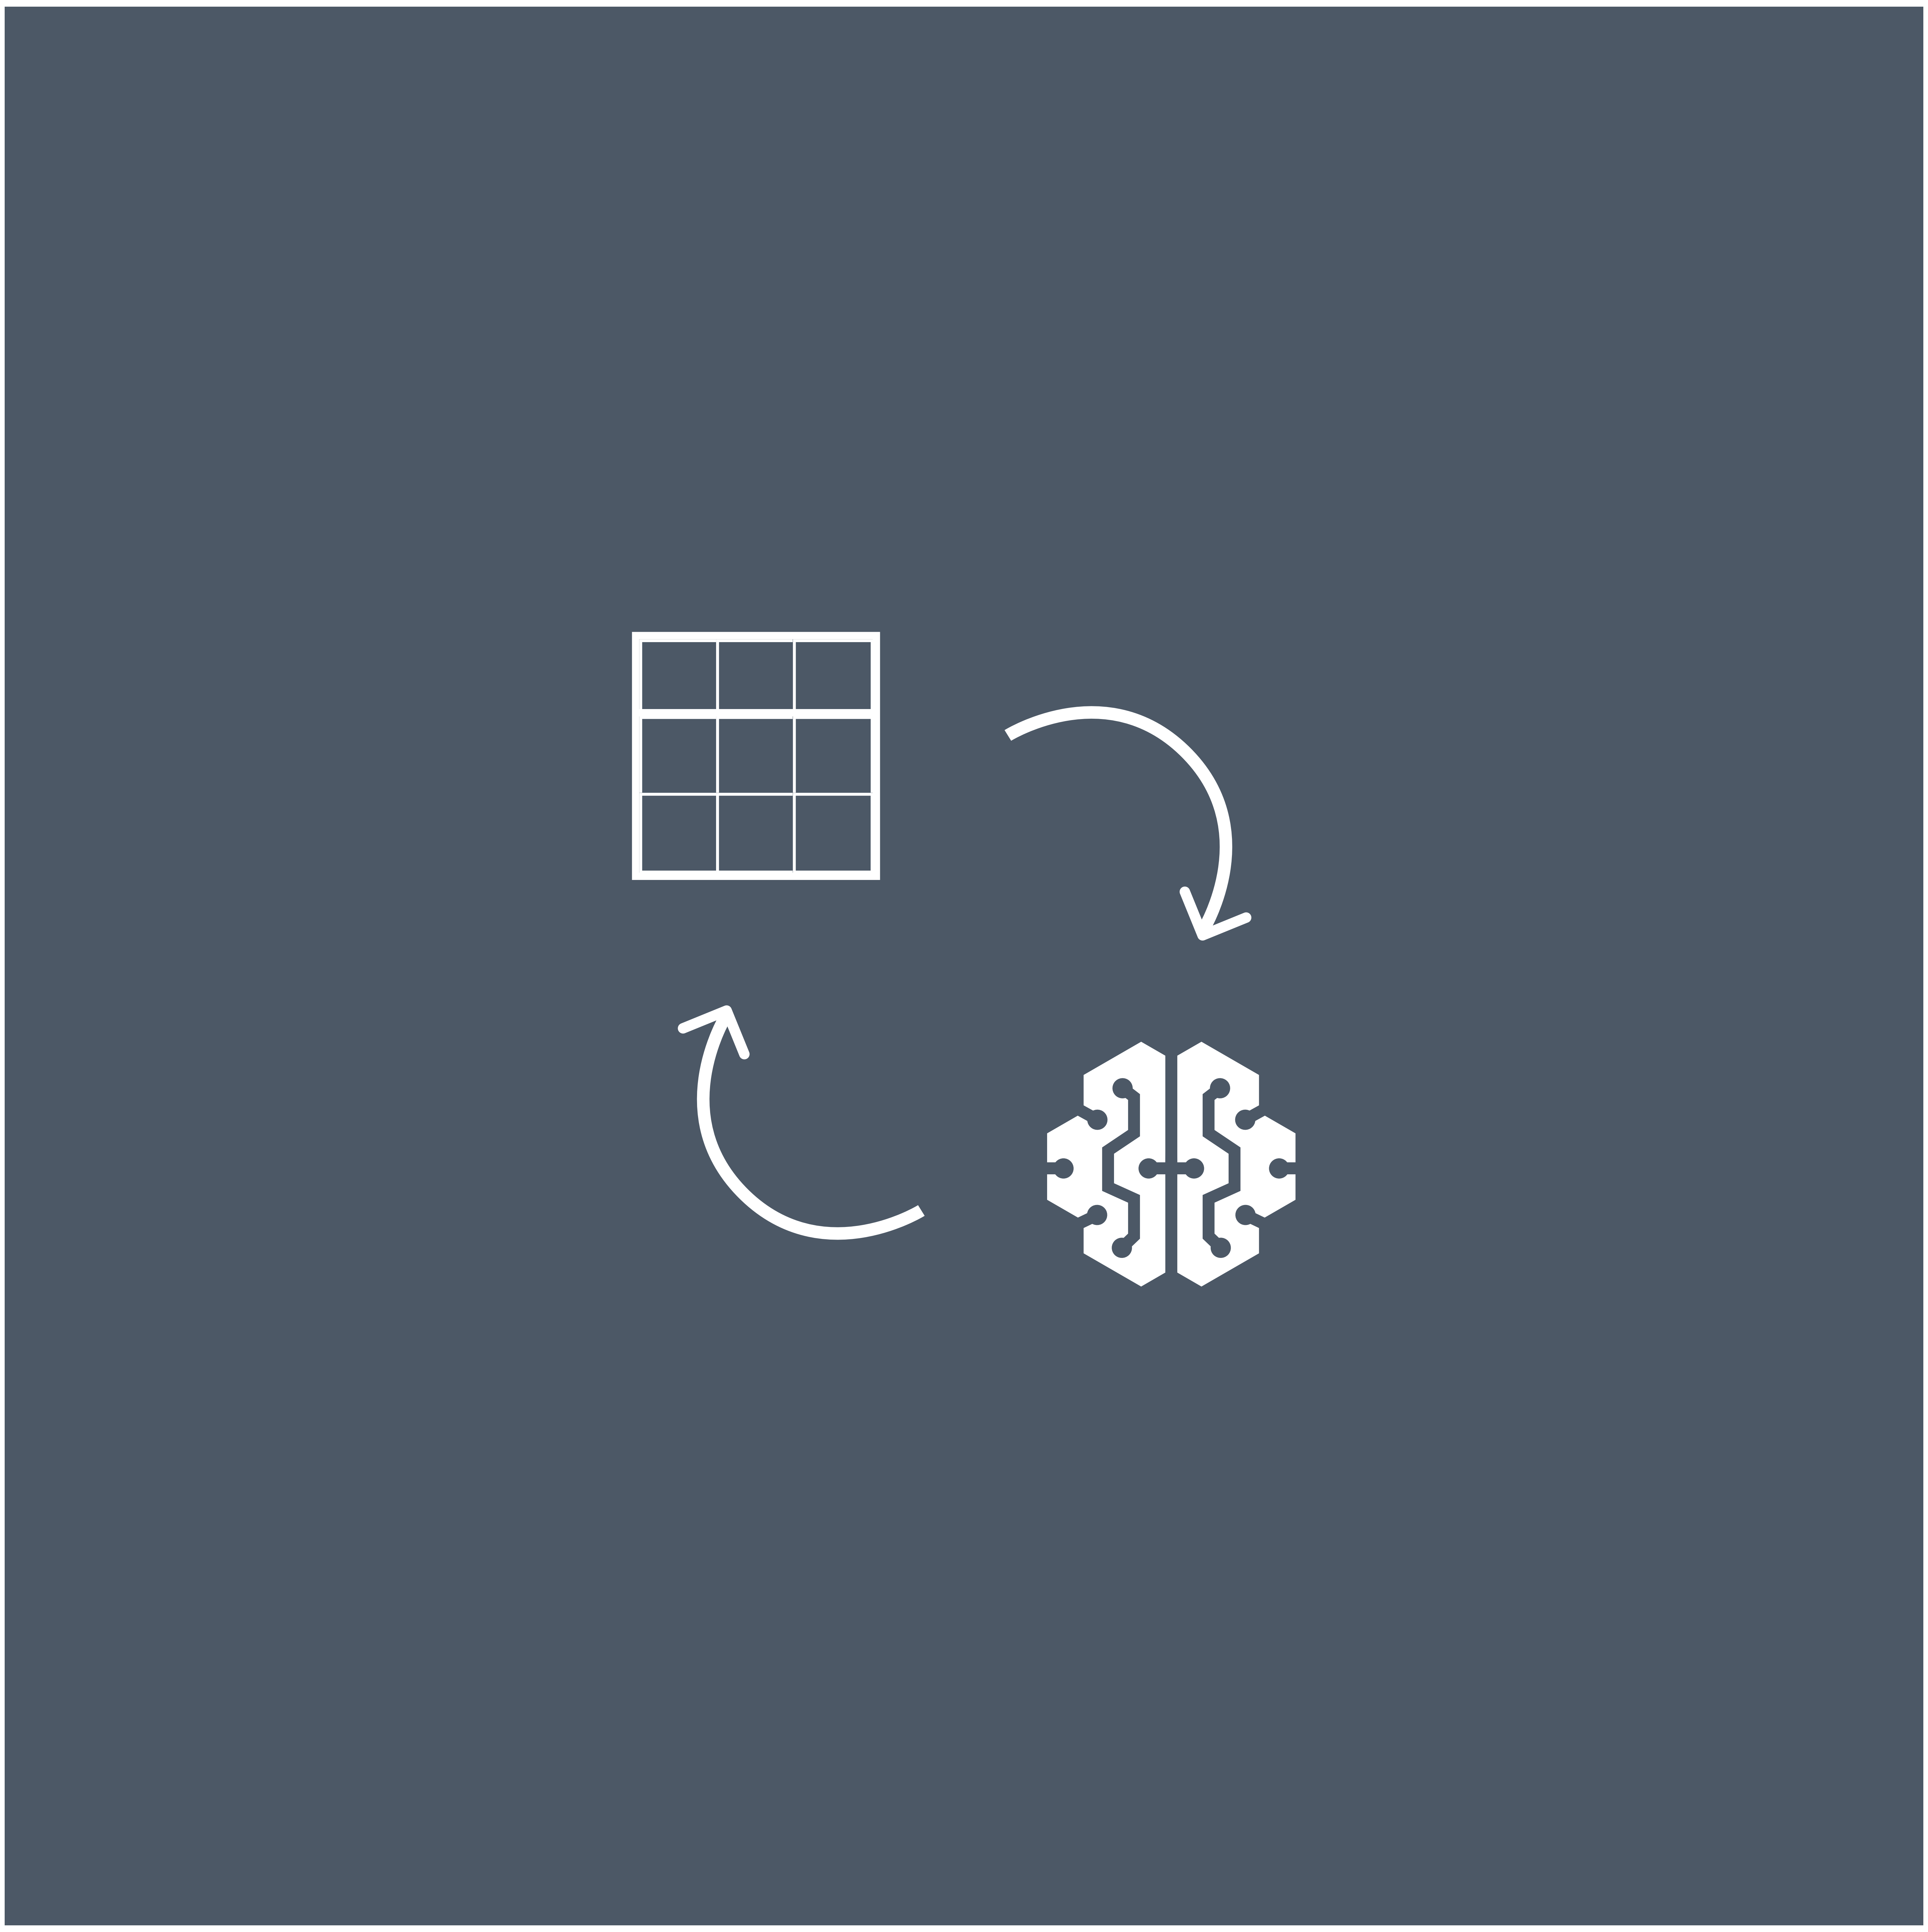

LEARNING

Data is to the neural network as experience is to humans.

A machine learning algorithm, in this case a neural network, uses data to find useful patterns and relationships. It uses these insights to learn and update itself.

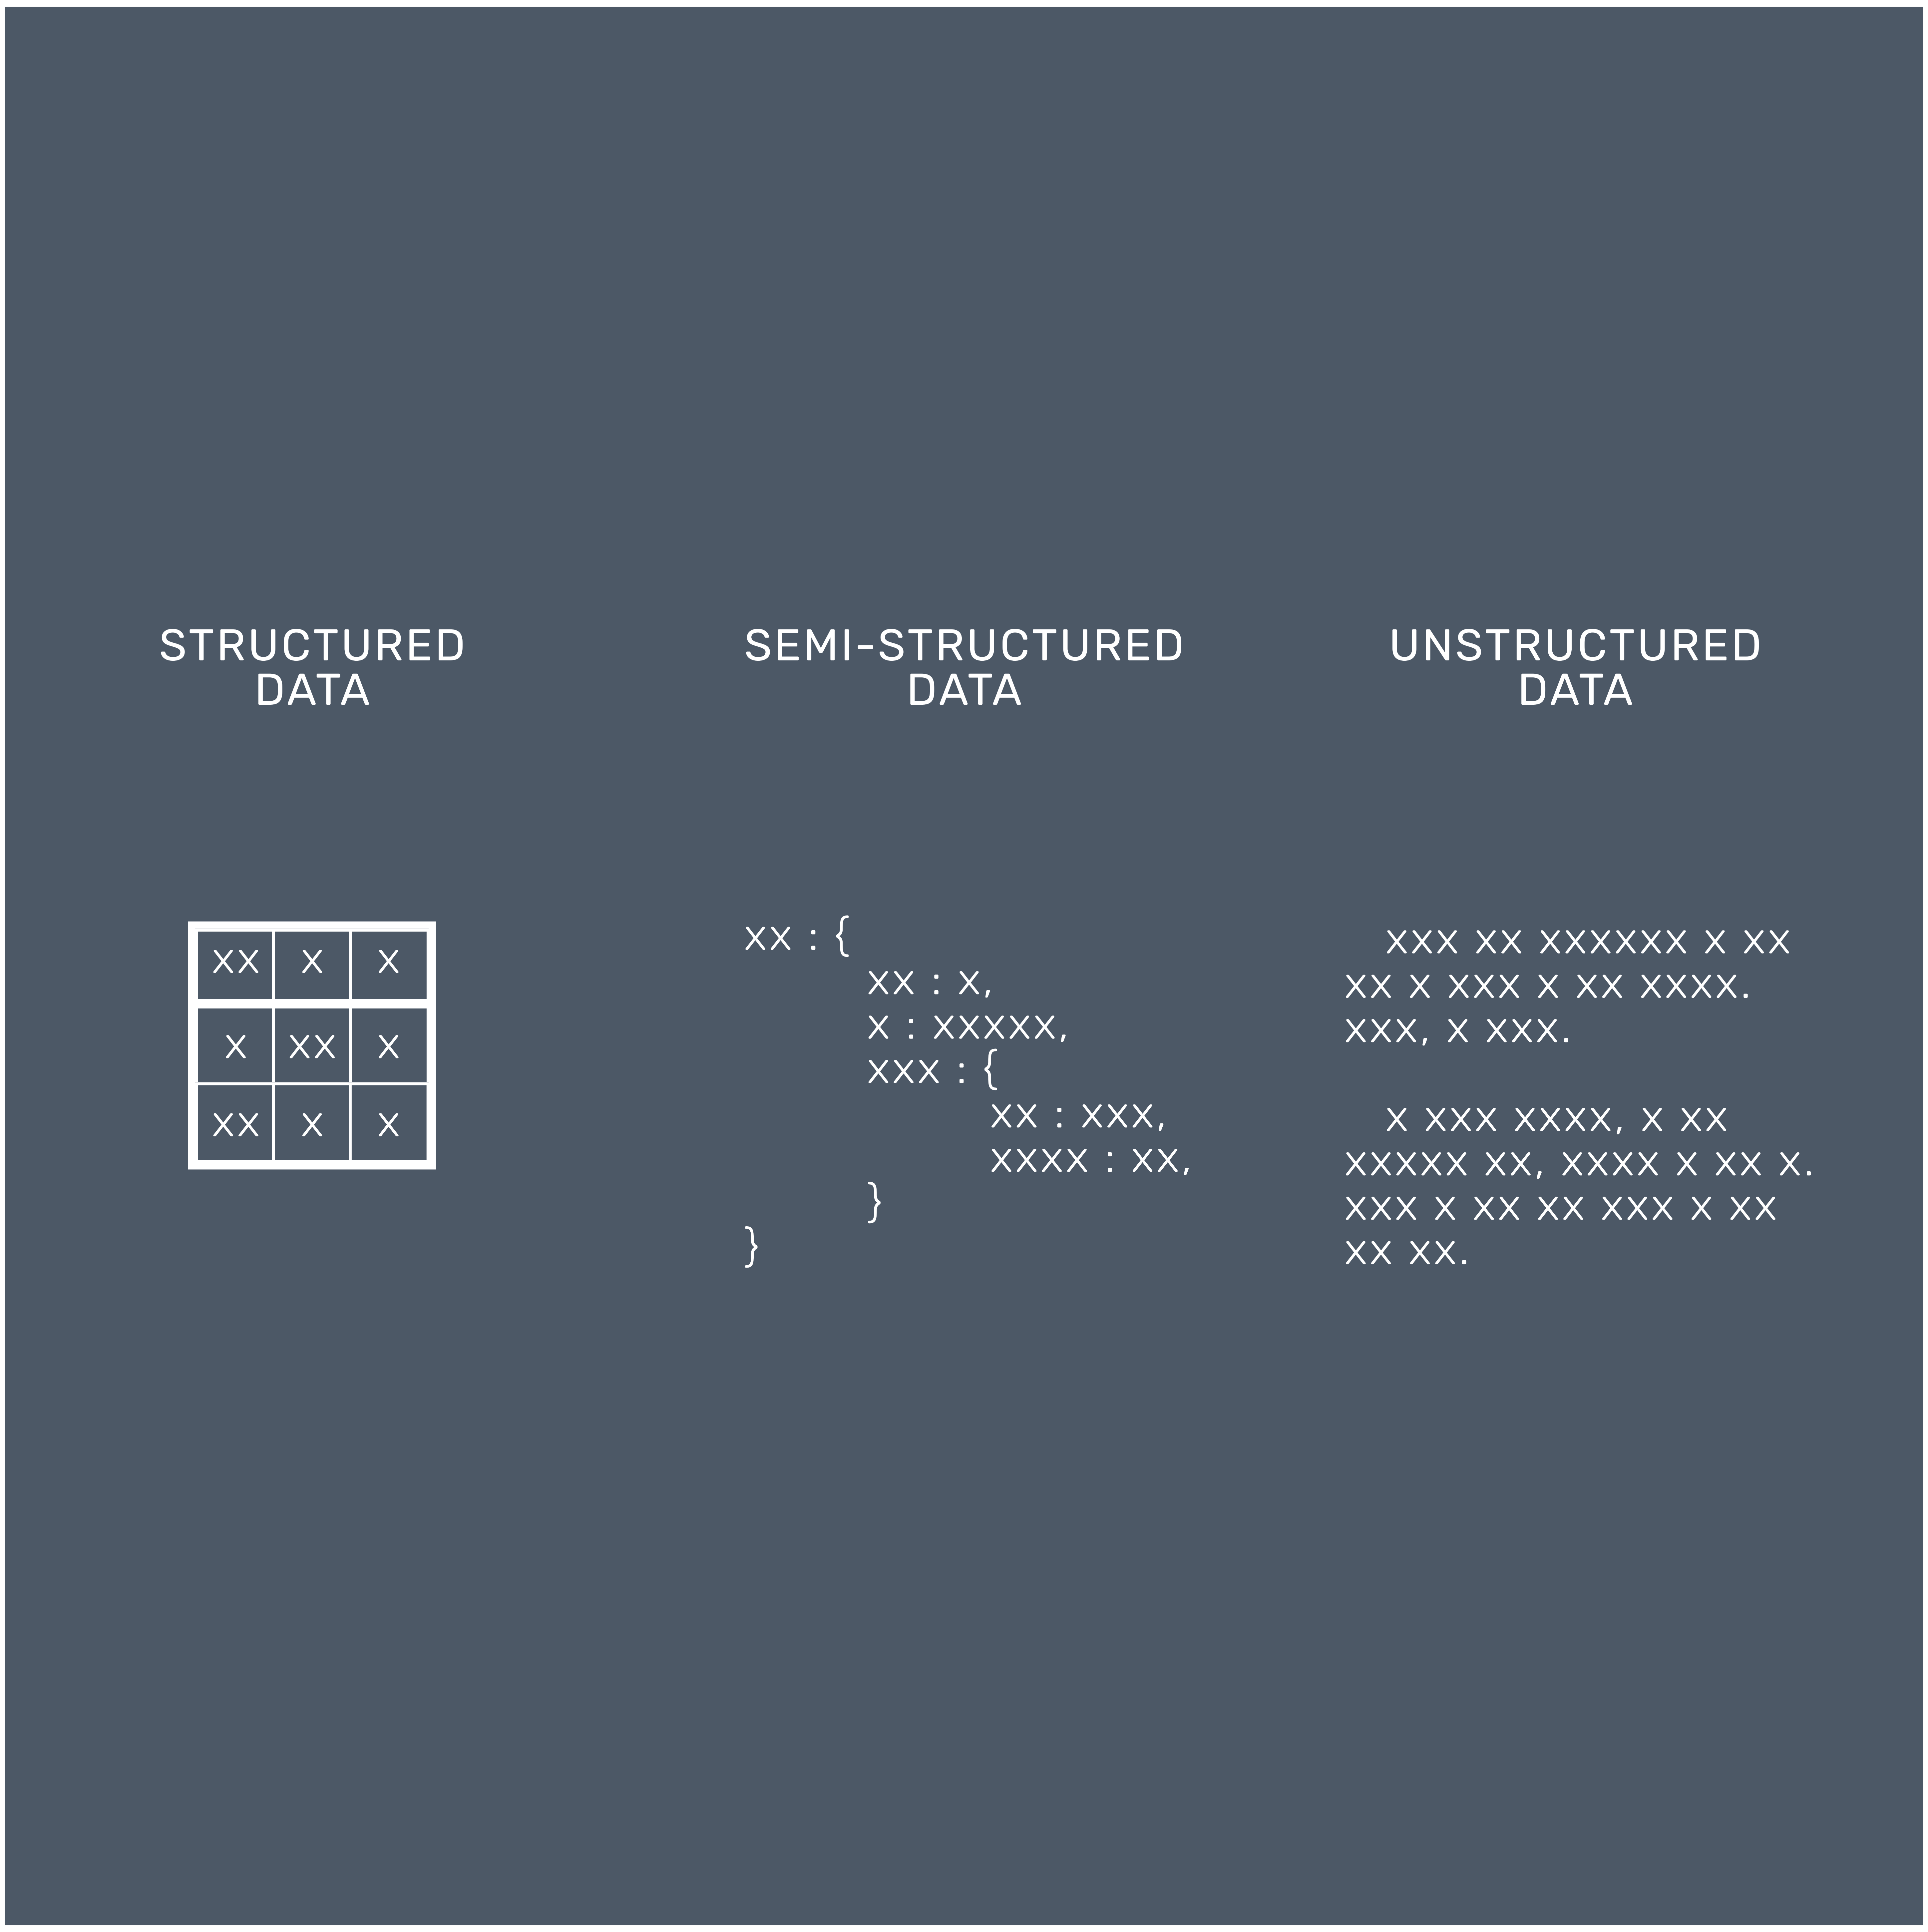

TYPES OF DATA

Data can come in many different forms. The most obvious form of data is the tabular format. This is an example of *structured* data, where each data point and its properties can be deciphered in a straightforward manner.

But data can come in other forms too.

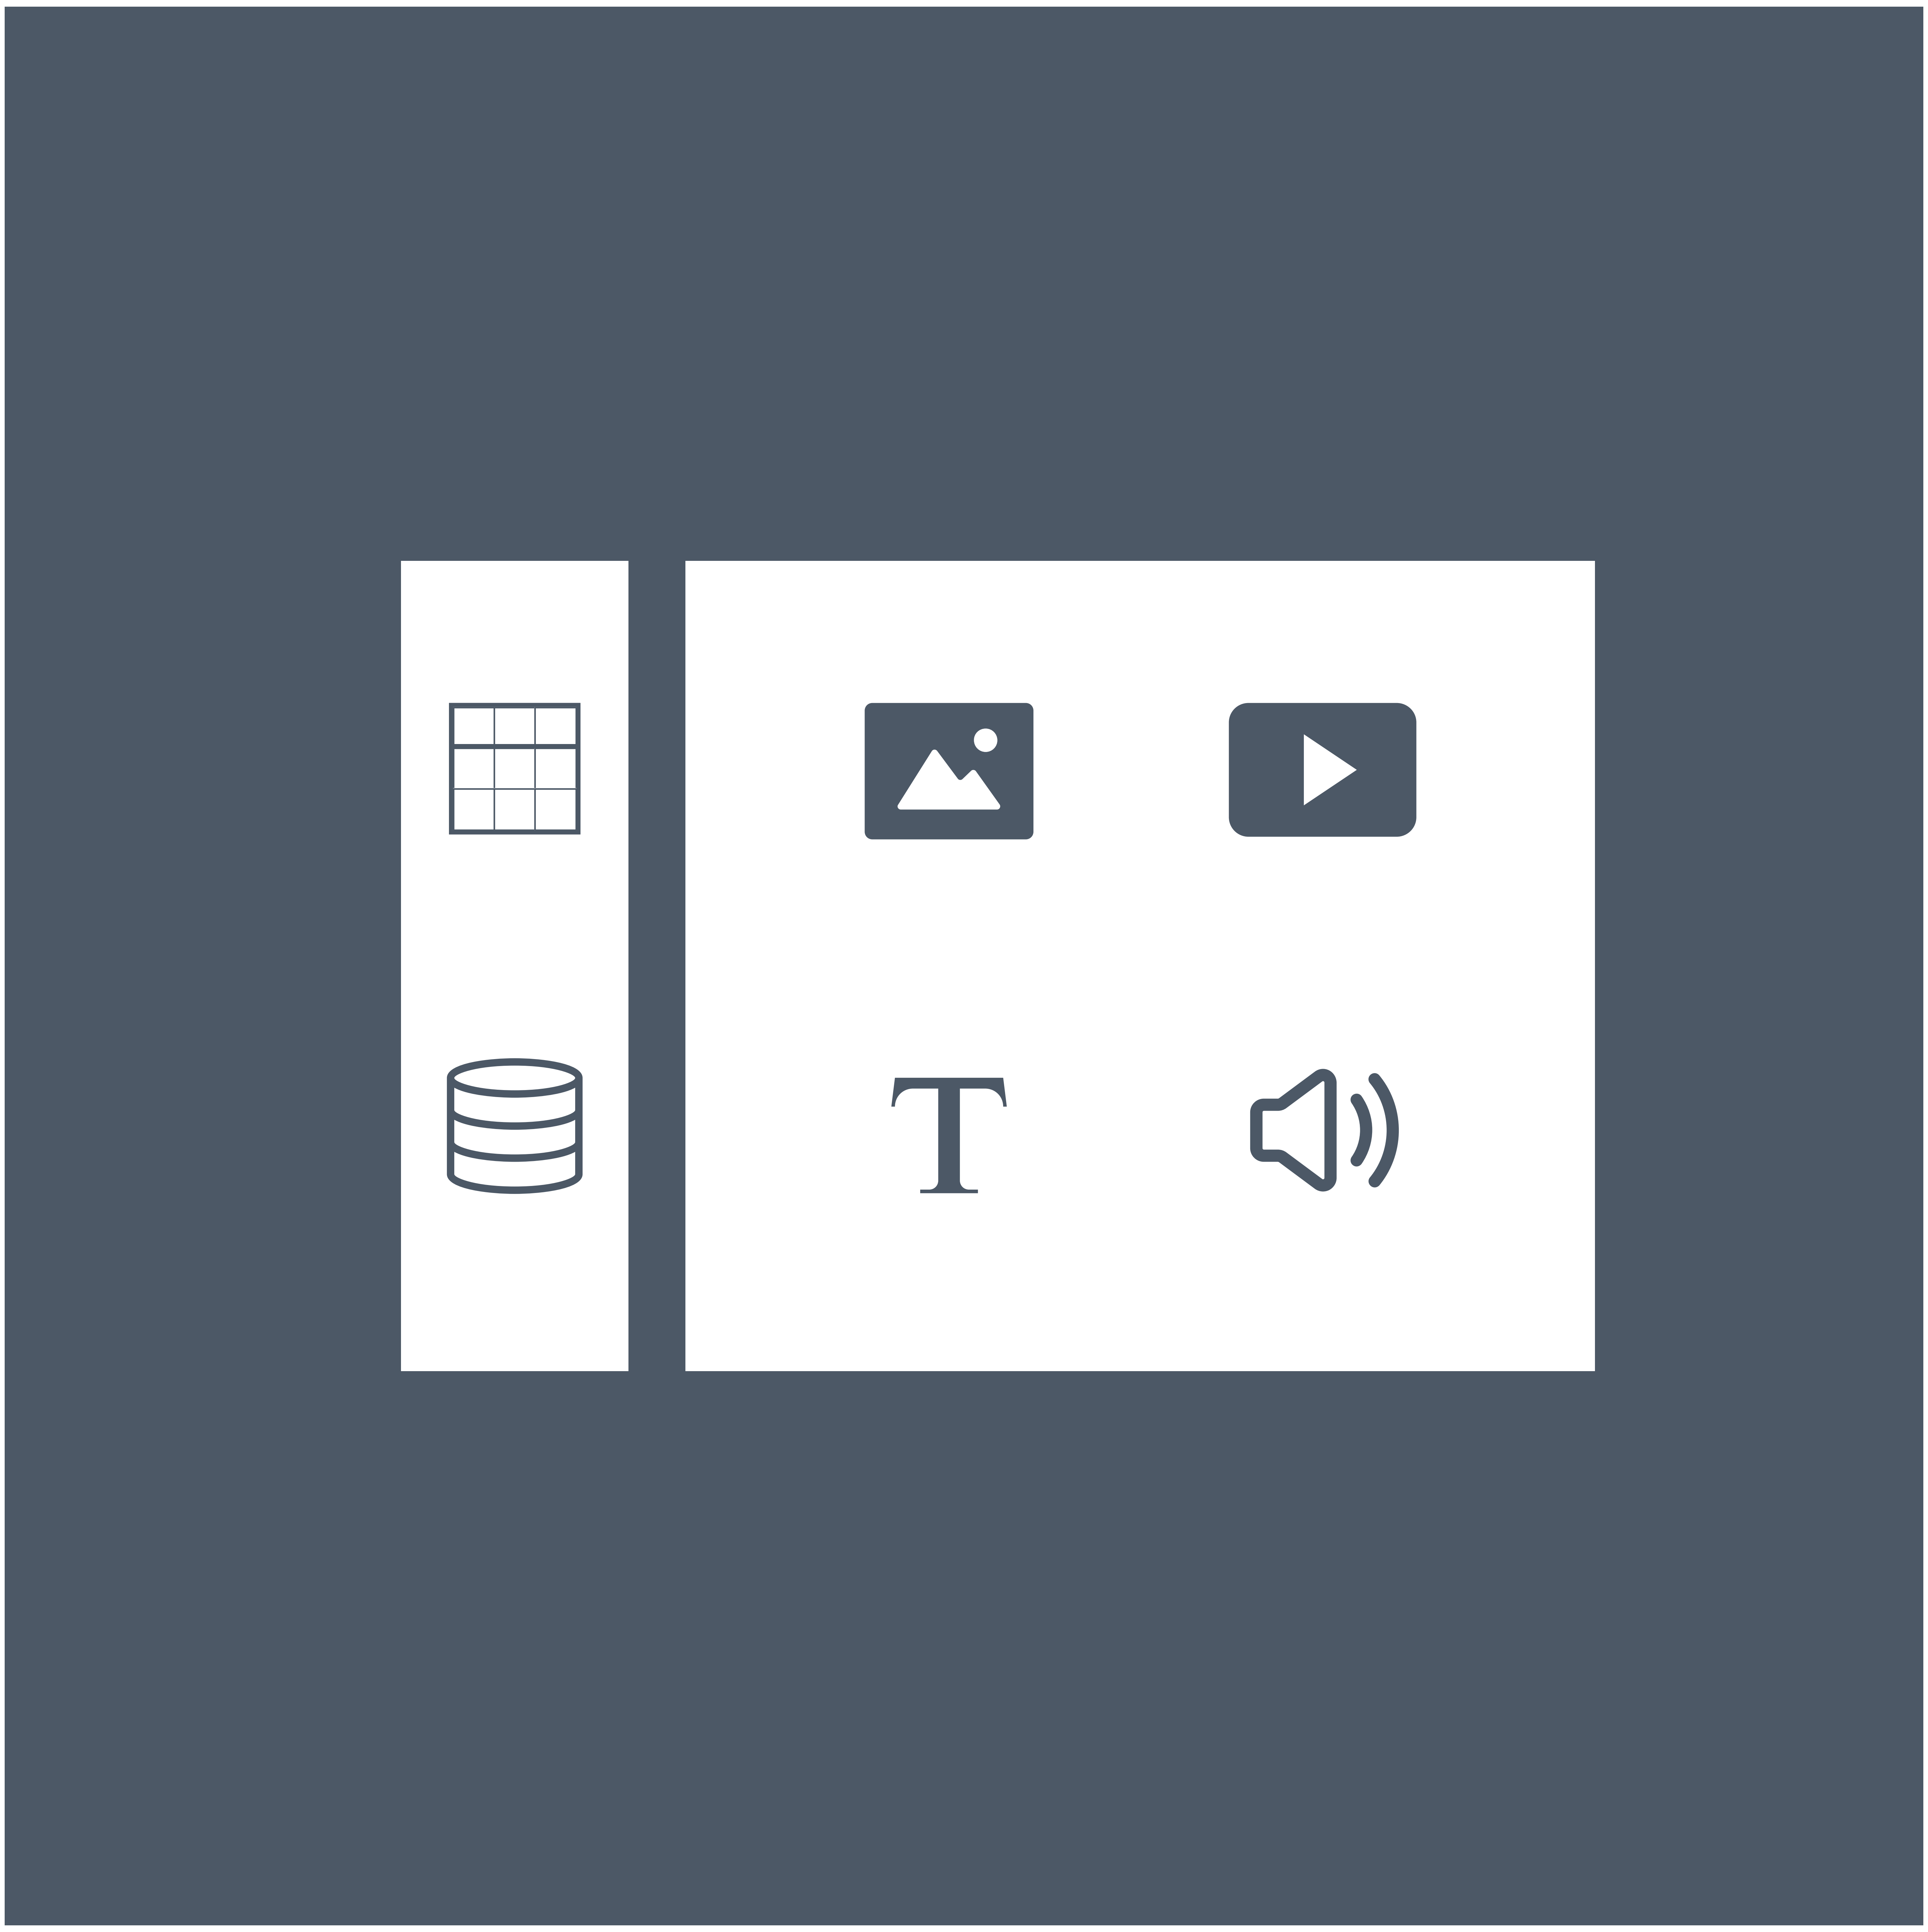

SOURCES OF DATA

In fact, most of the data around us are in the *unstructured* form. According to projections from IDC, 80 percent of worldwide data will be unstructured by 2025.

And indeed, most of the exciting innovations in deep learning today come from unstructured data, such as text, images, videos, and so on.

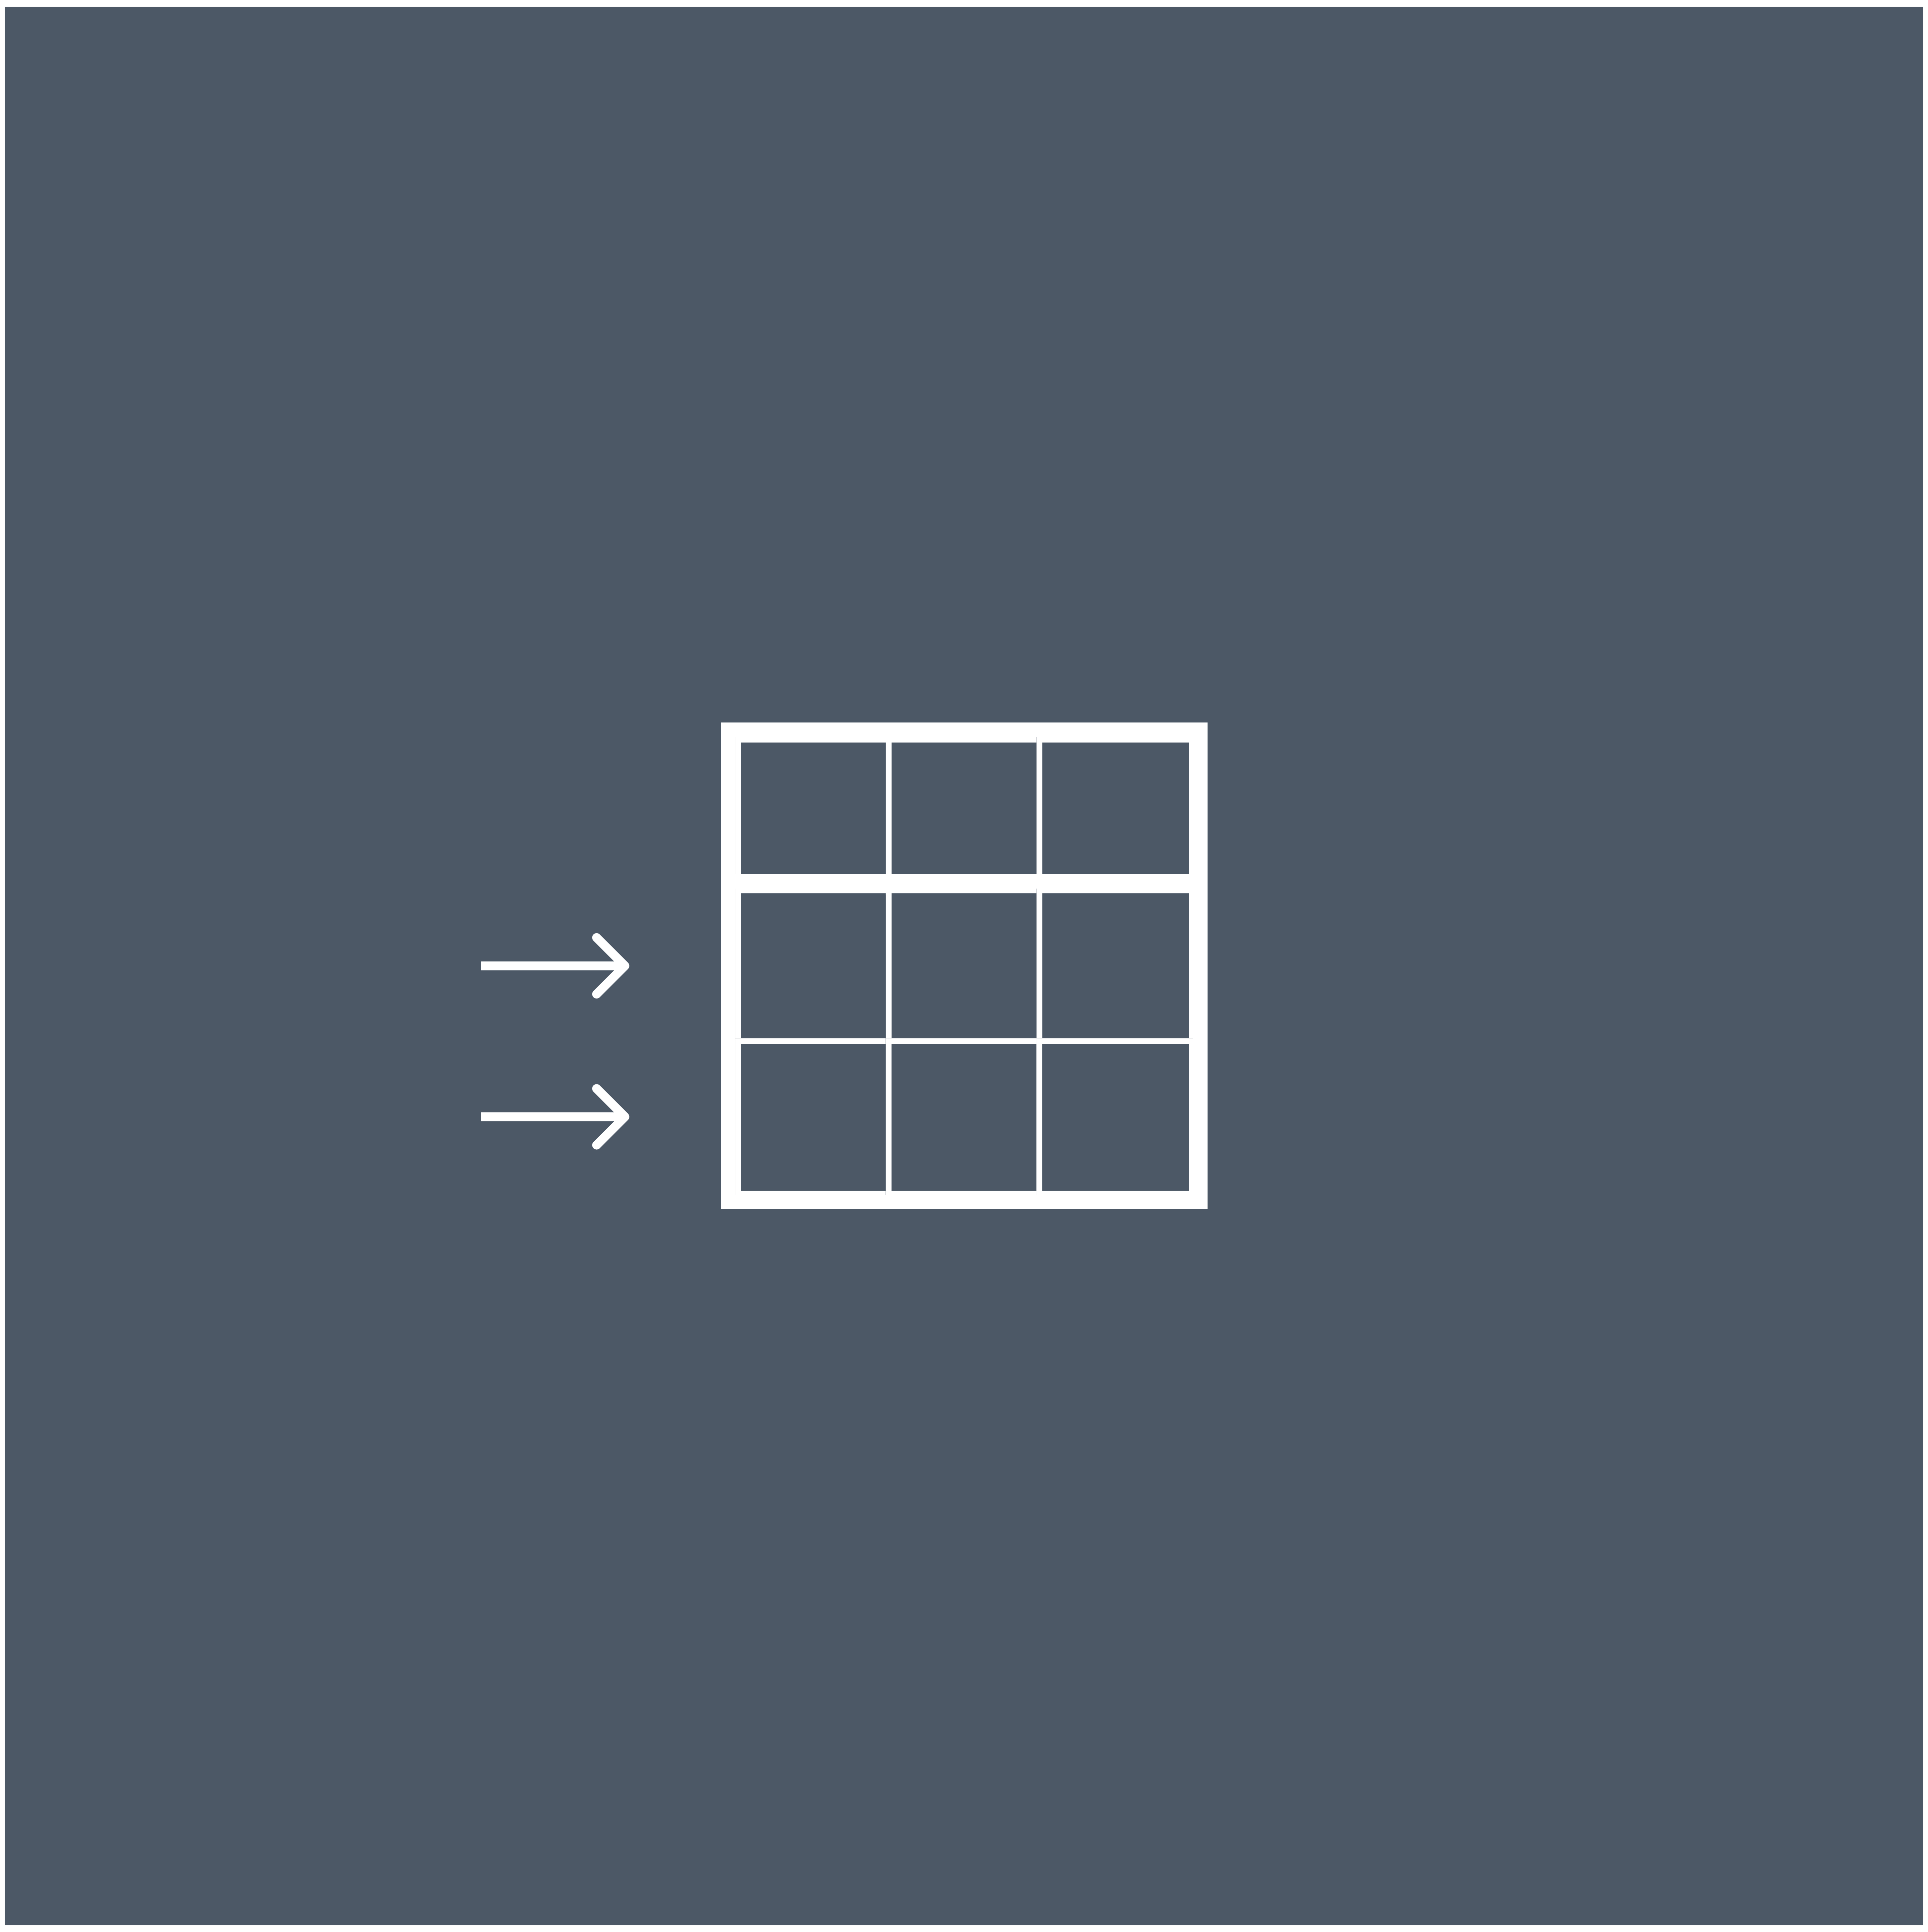

A DATASET

Now let's look at how we should prepare a *dataset* for the neural network.

A dataset is composed of many data points. A *data point* is a single observation that we collect and record.

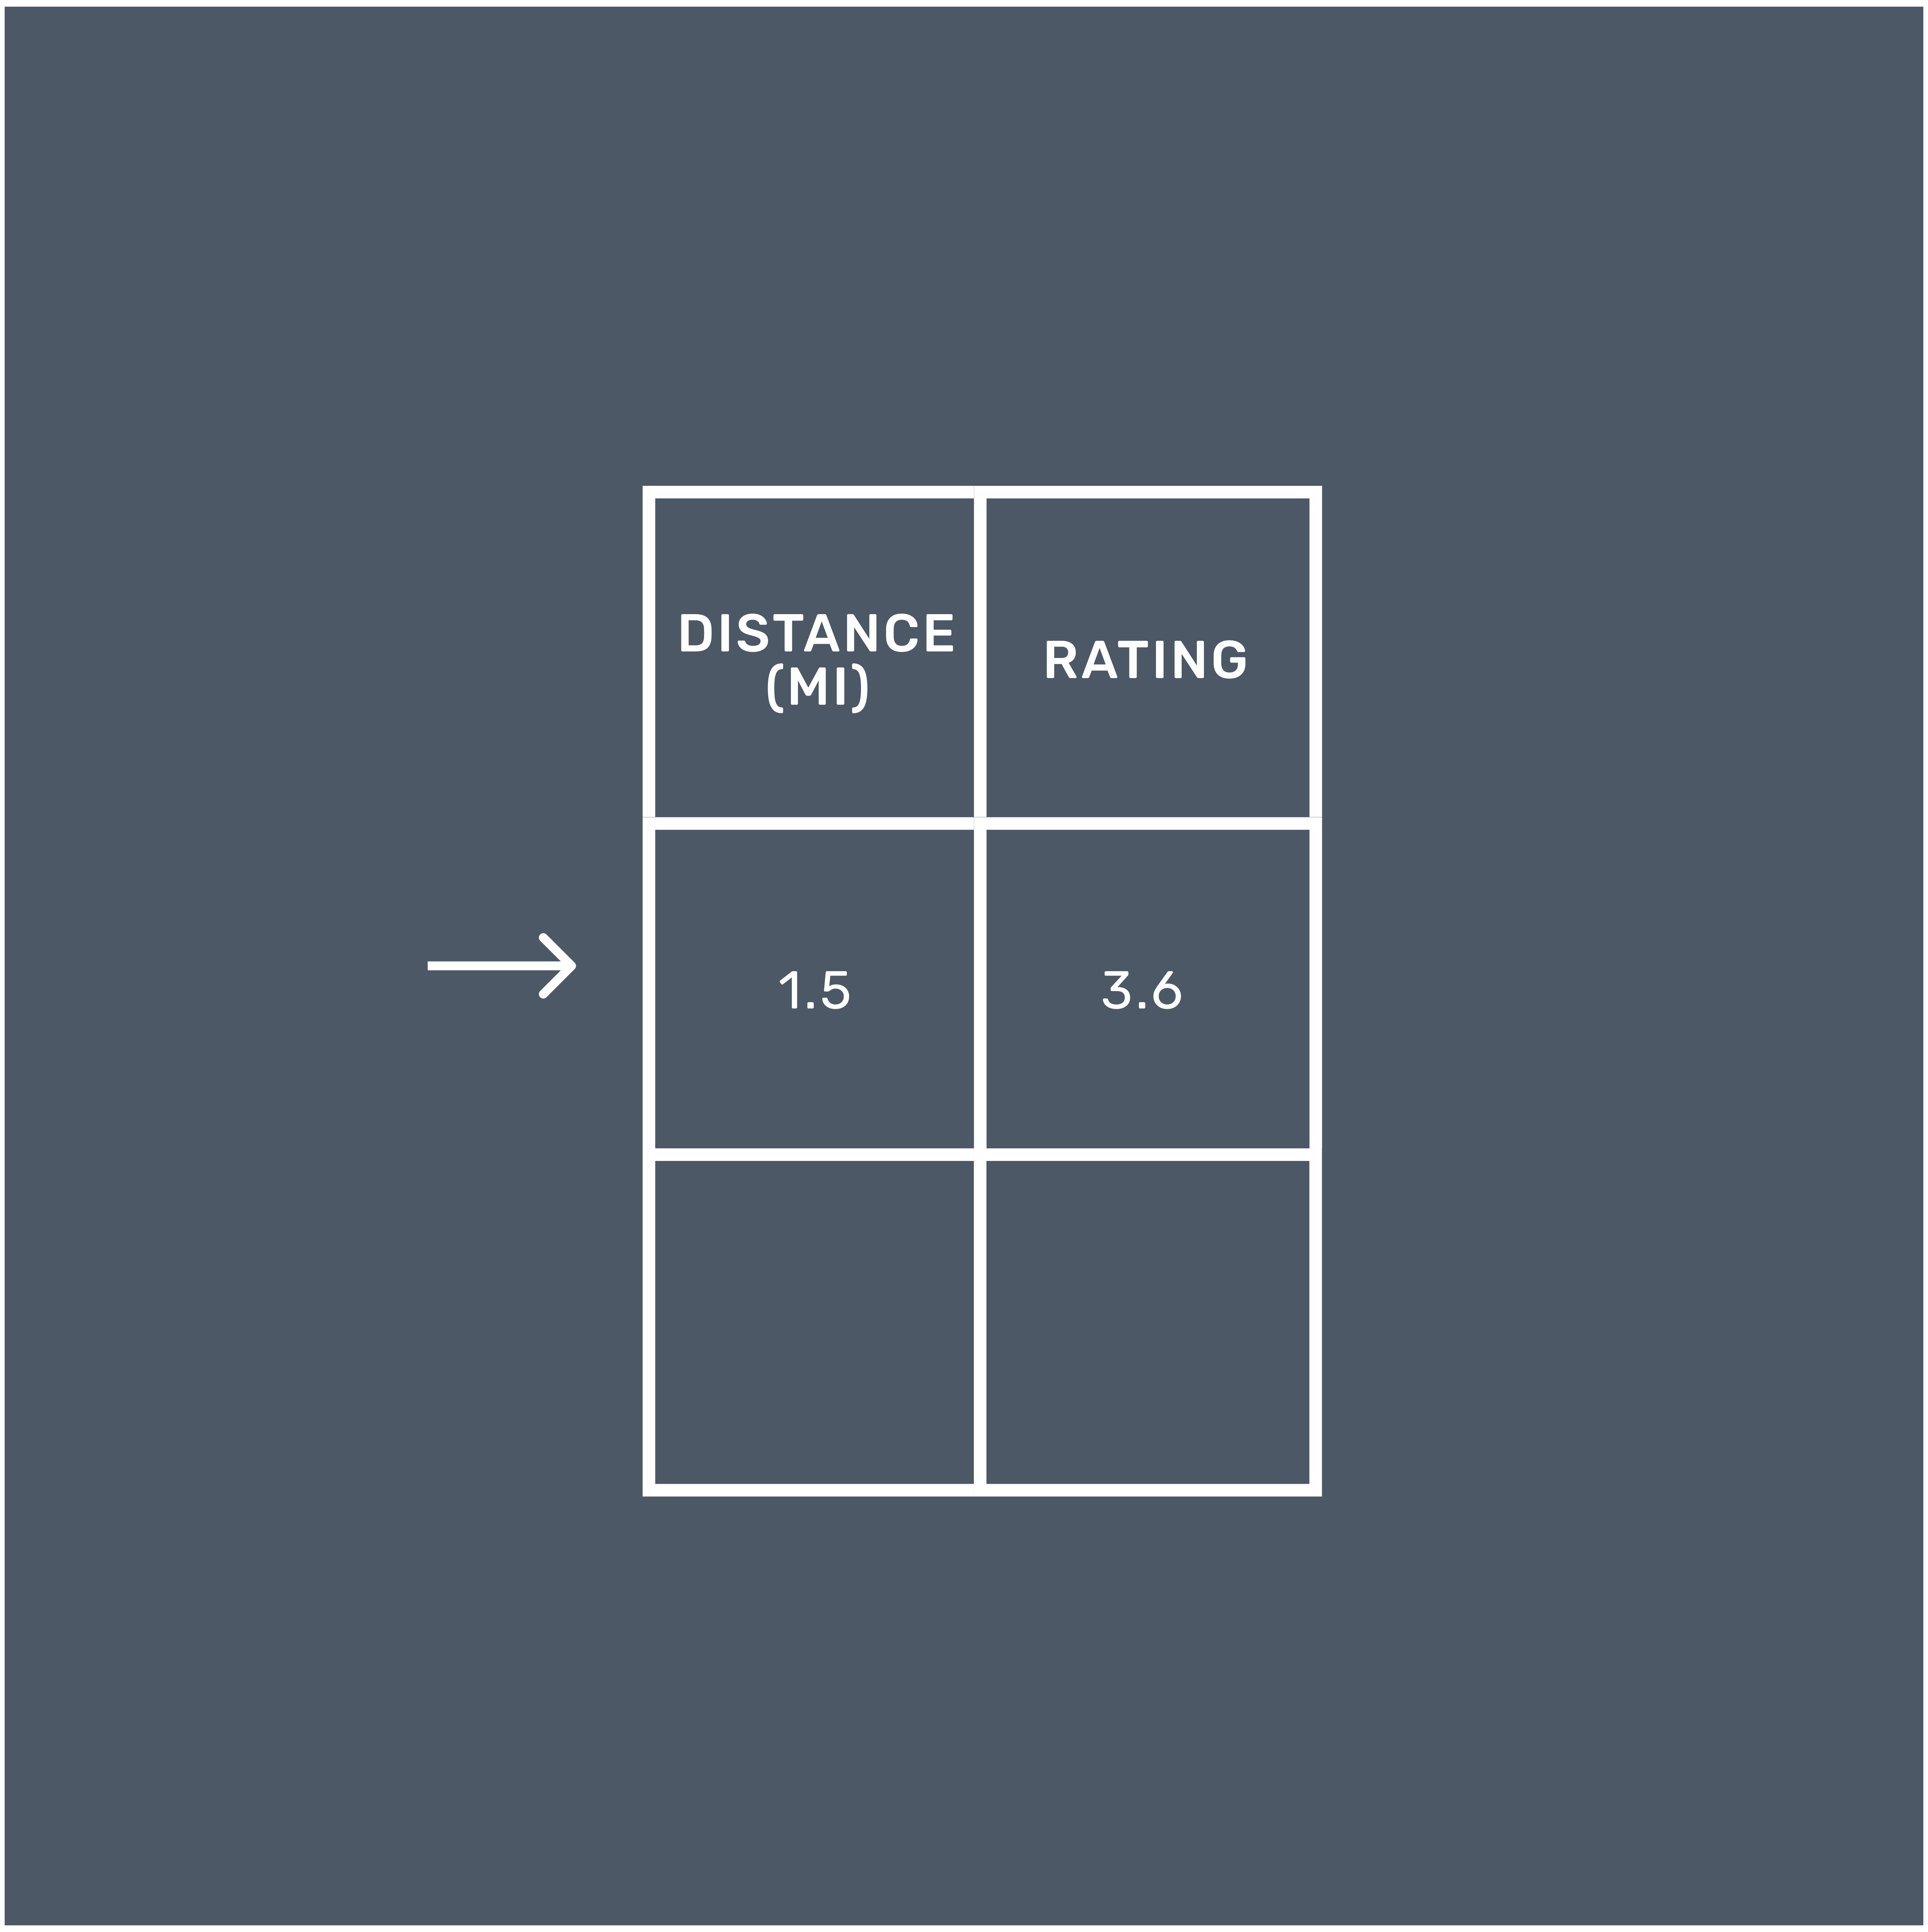

EXAMPLE

Let's take the example of hotel room rates, a dataset we'll use throughout this book.

Each data point represents a hotel. Here we have a hotel with a distance of 1.5 miles from the city center and a guest rating of 3.6 stars.

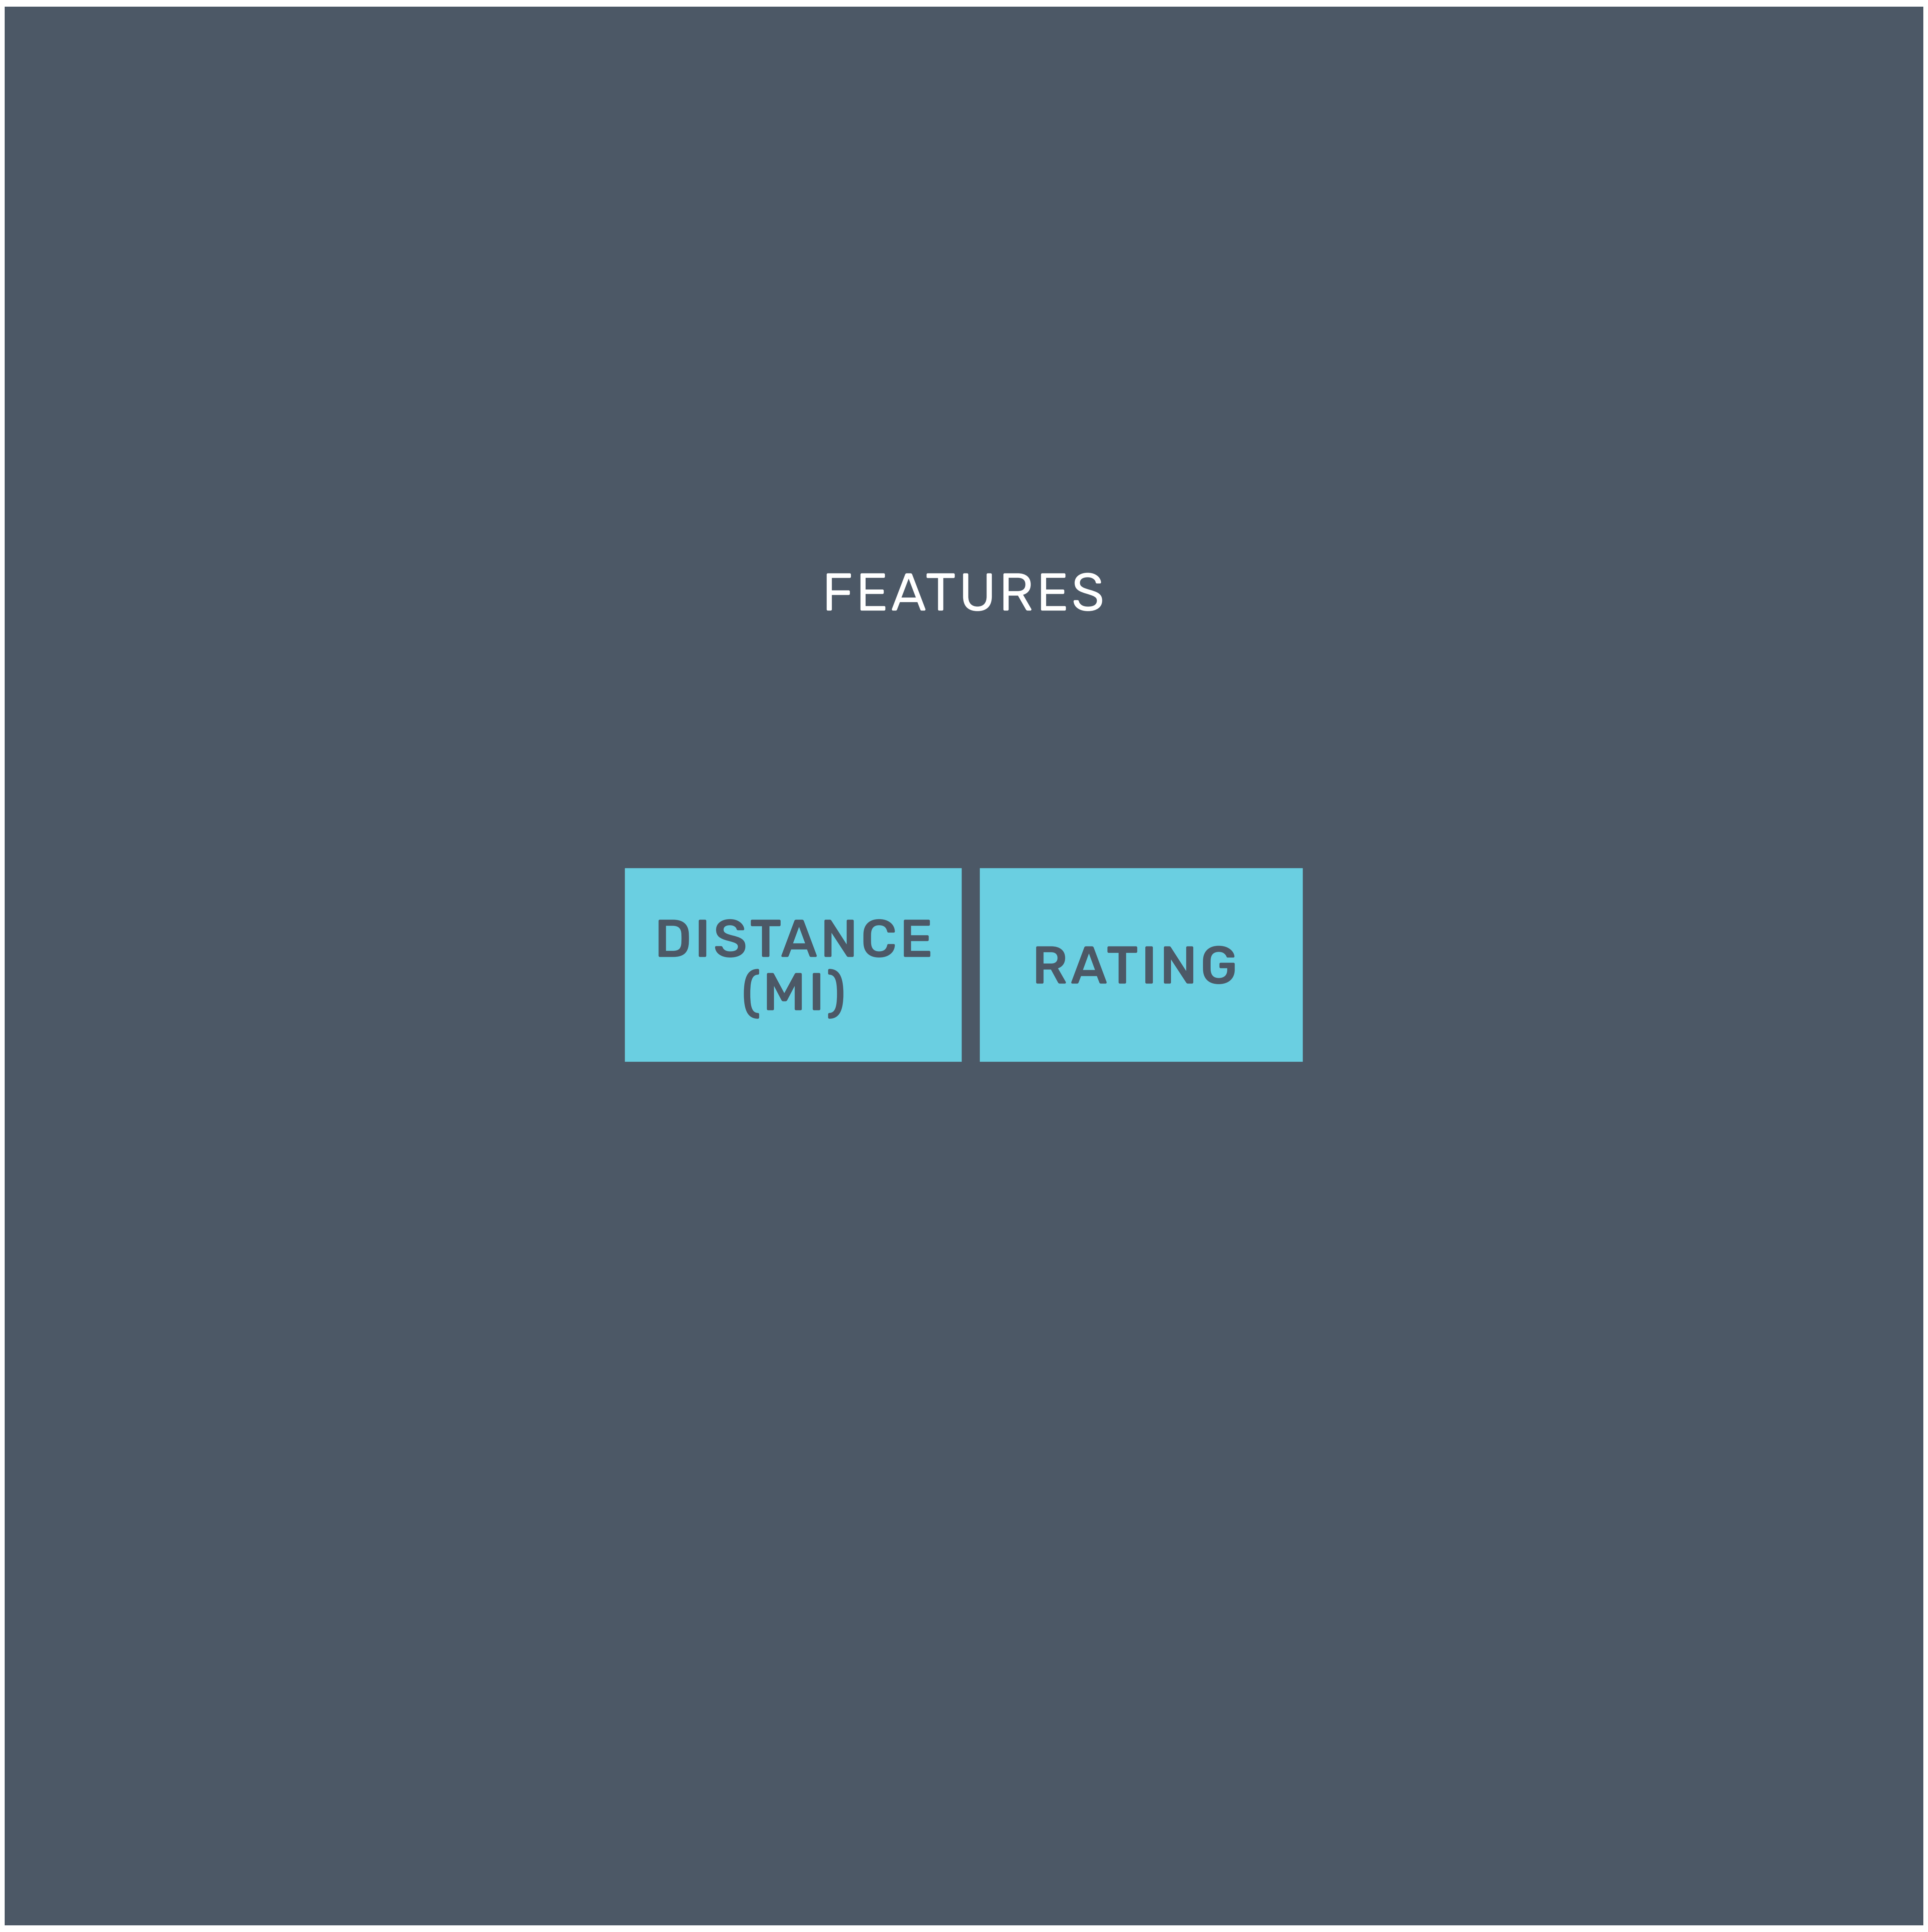

FEATURES

These two pieces of information are called *features*. Features describe the properties of a data point.

Each data point in a dataset is described with the same features, of course with different values, making each of them unique.

From now on, we'll refer to these two features as *distance* and *rating* for short.

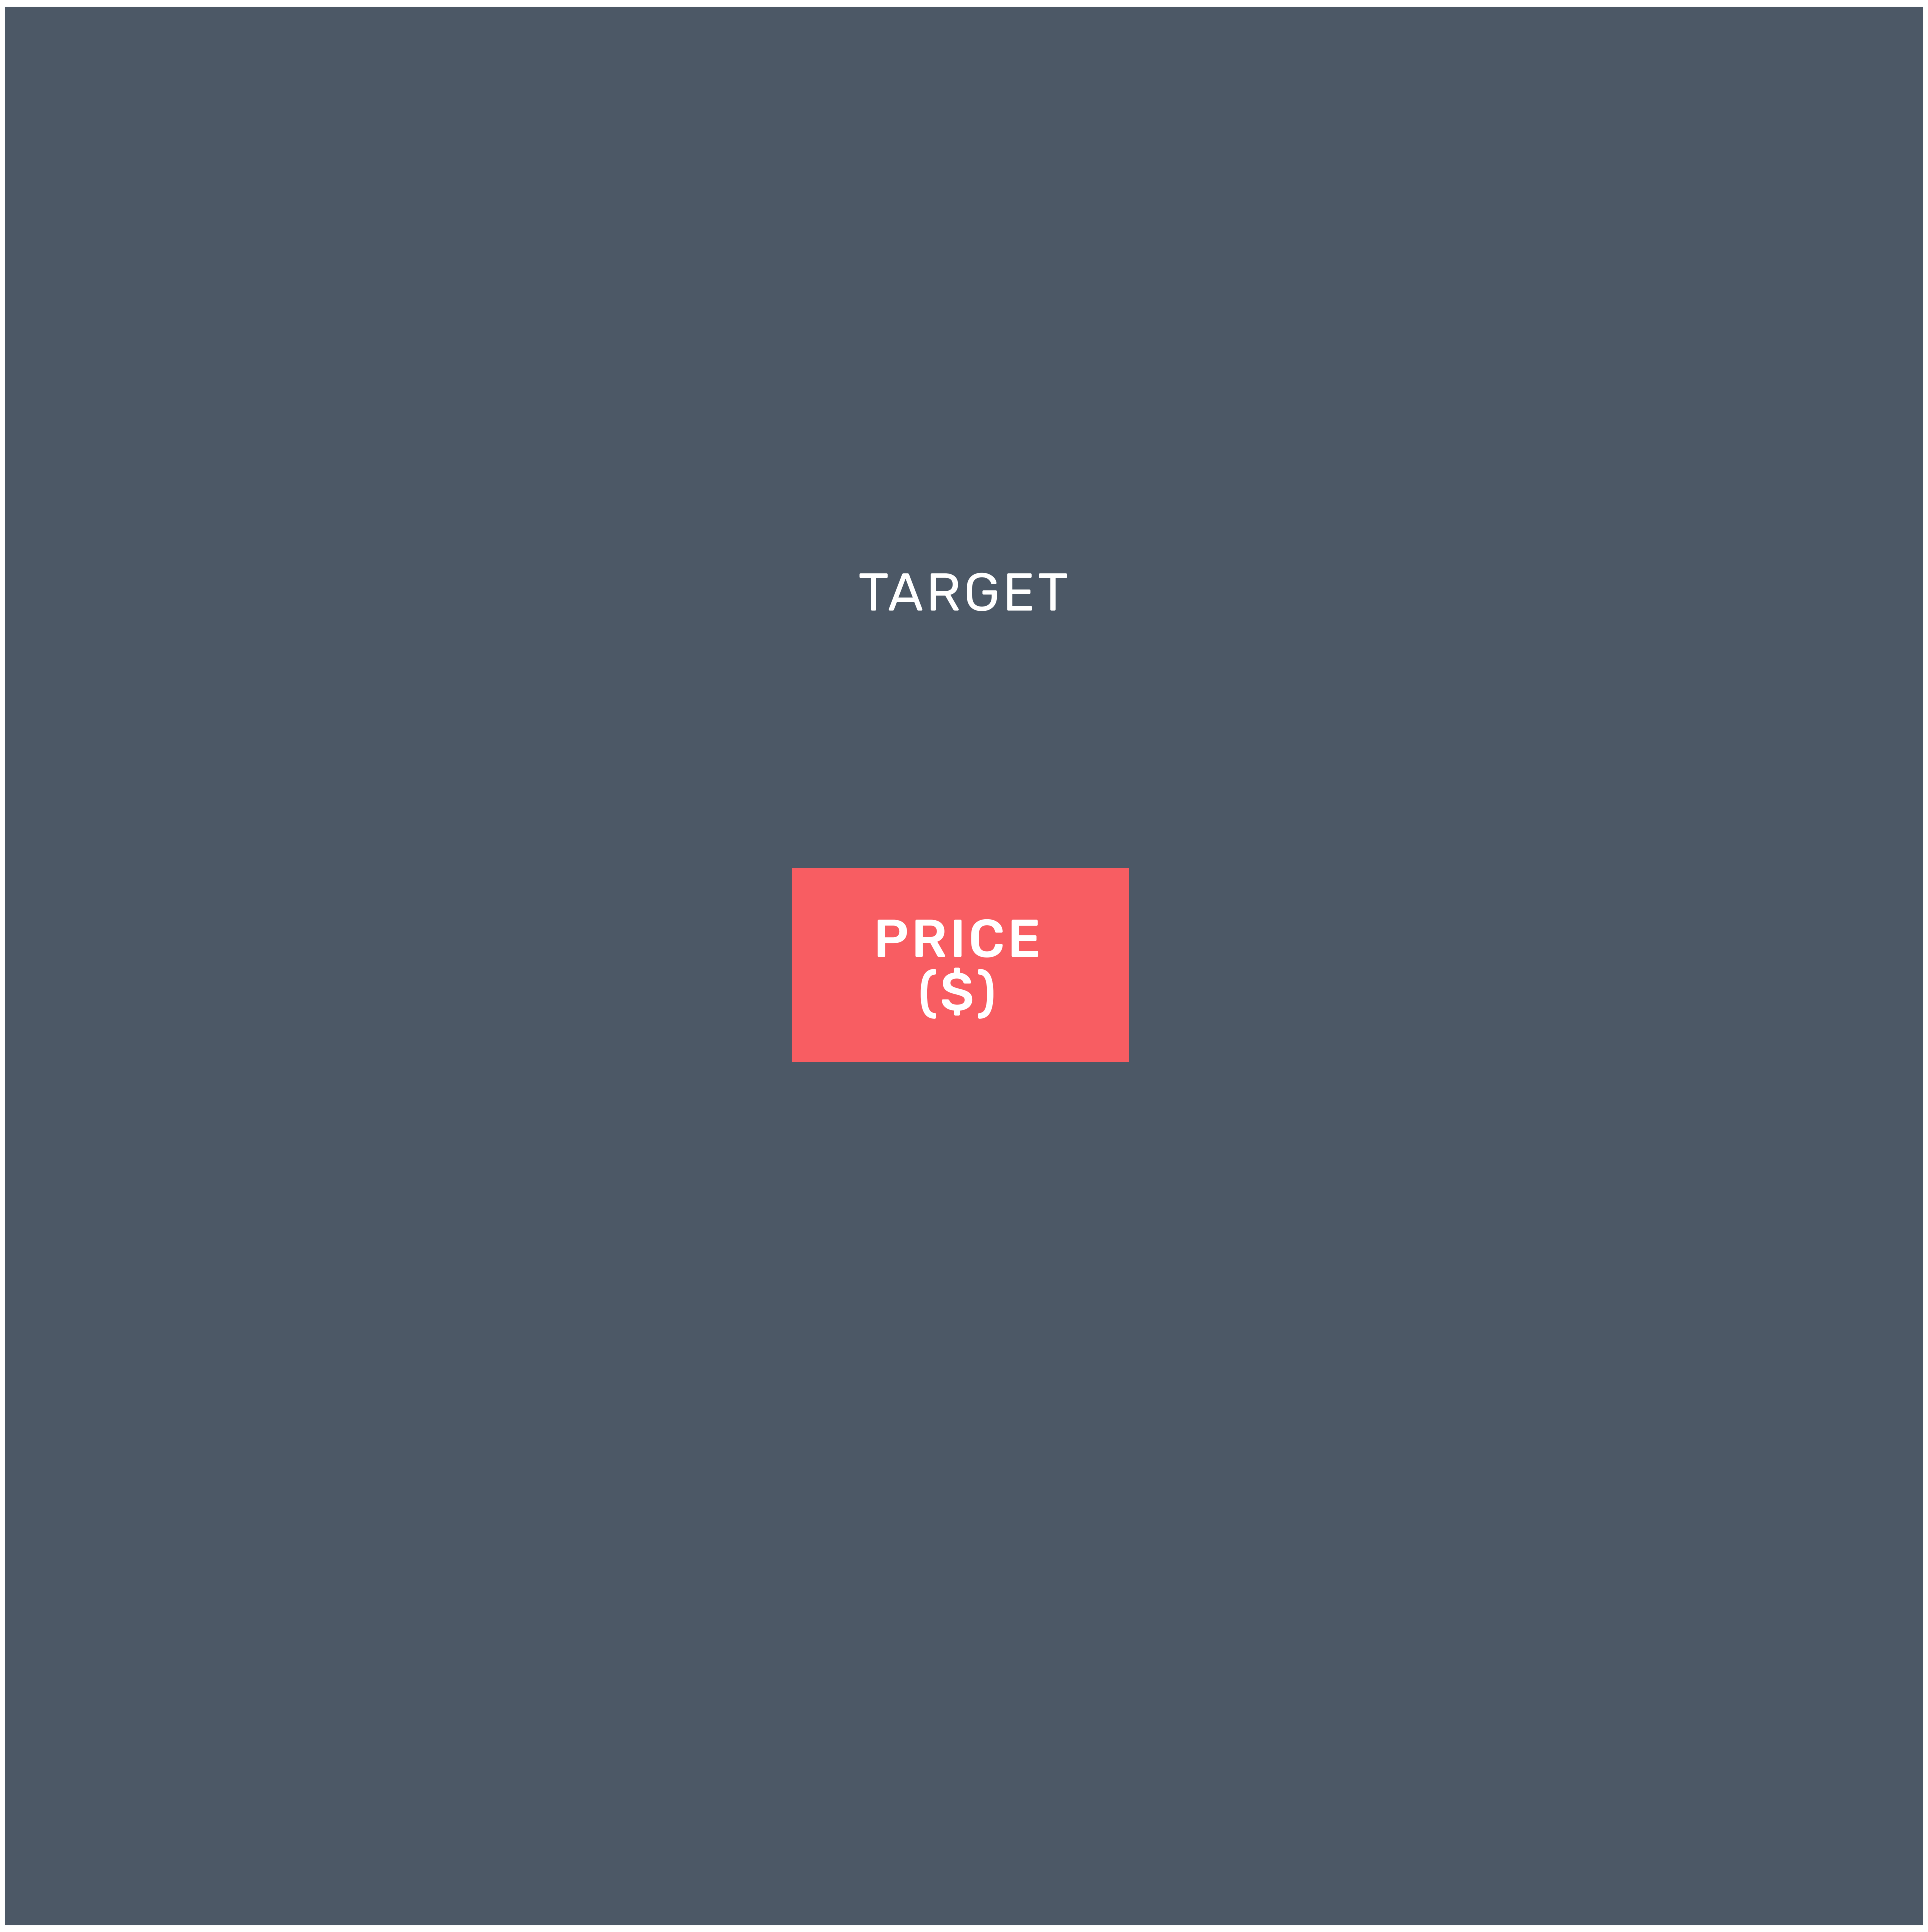

TARGET

Recall that the goal of a neural network is to make predictions.

In this example, we want to predict the average daily room rate (or *price* for short) for any given hotel. This means, given the two features earlier, we want to predict how much each hotel will cost.

The is called the *target*. In other resources, you may also find the term *label* being used.

DIST (MI)	RATING	PRICE ($)
0.8	2.7	147
1.5	3.6	136
...	...	...
19.4	4.8	209

TRAINING

We'll give the neural network enough data points containing the features and target values, which it will learn from.

A machine learning task where we specify the target is called *supervised learning*, which will be our focus in this book.

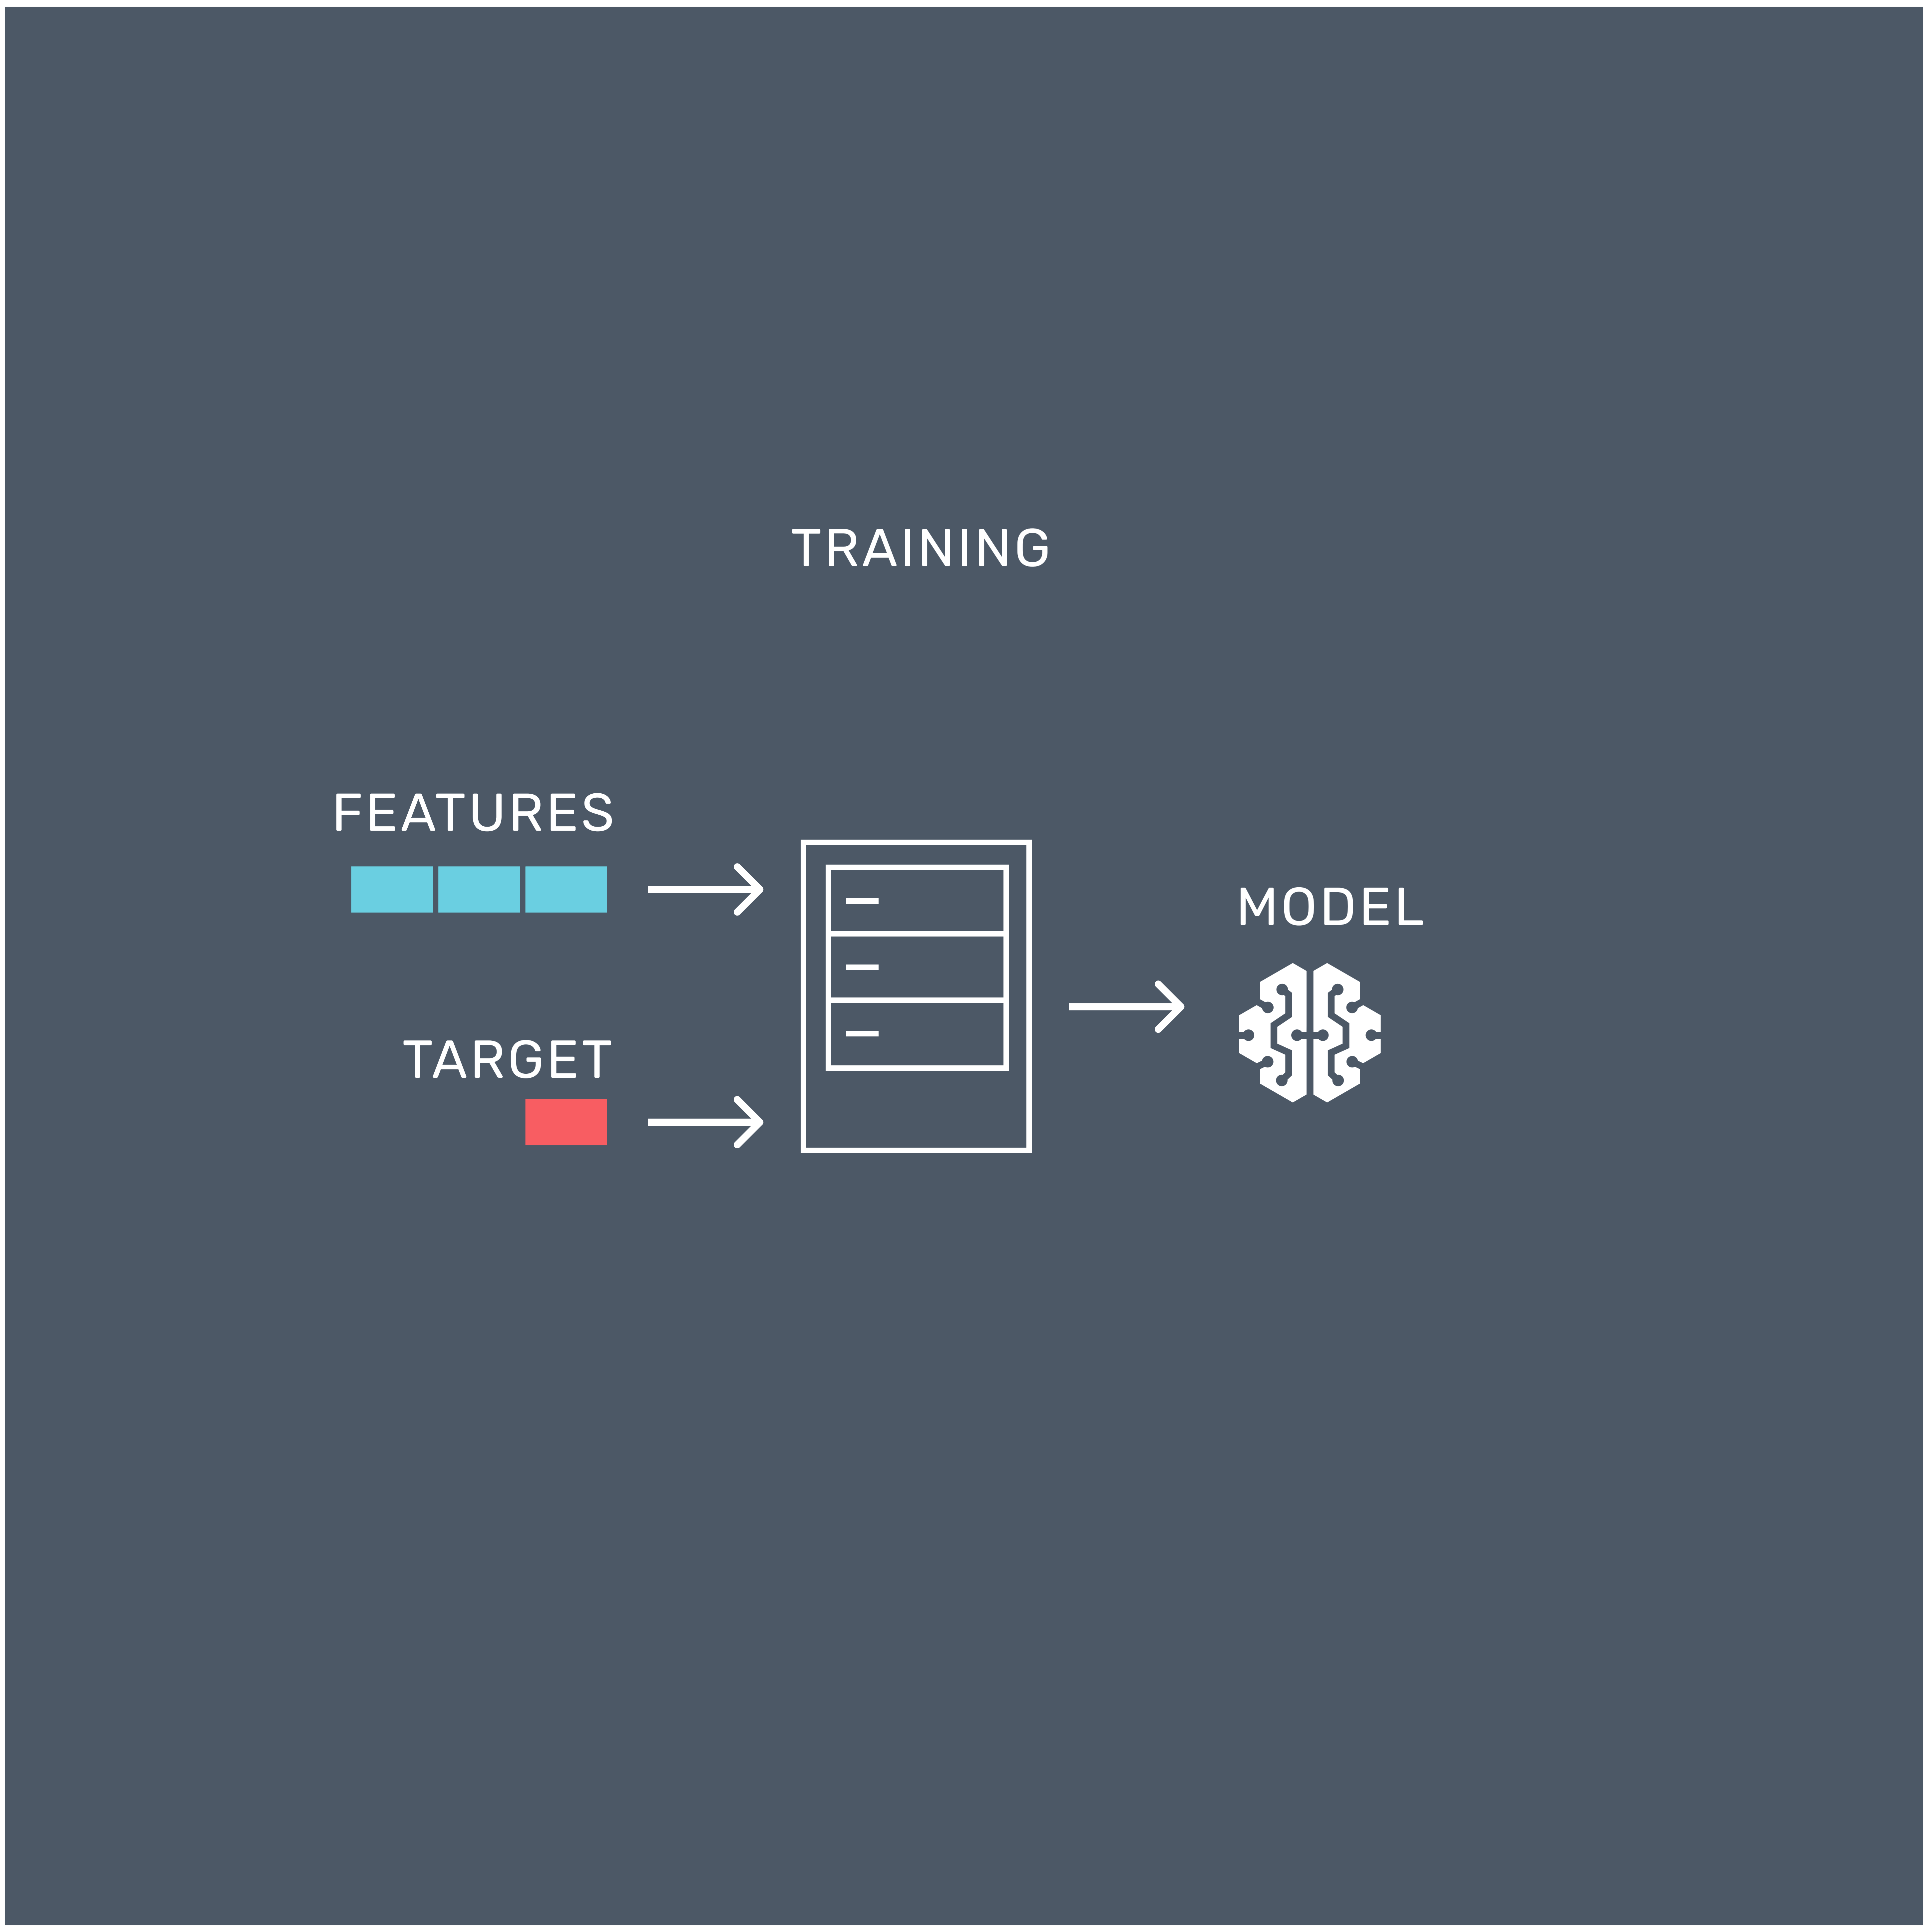

TRAINING

We have just described a process called *training*.

During training, the neural network learns from the data and updates its parameters. By this point, we'll have a trained *model*.

In short, given a dataset containing features and target, we get a model.

That is why the training process is sometimes also called ‘fitting the model to the data’.

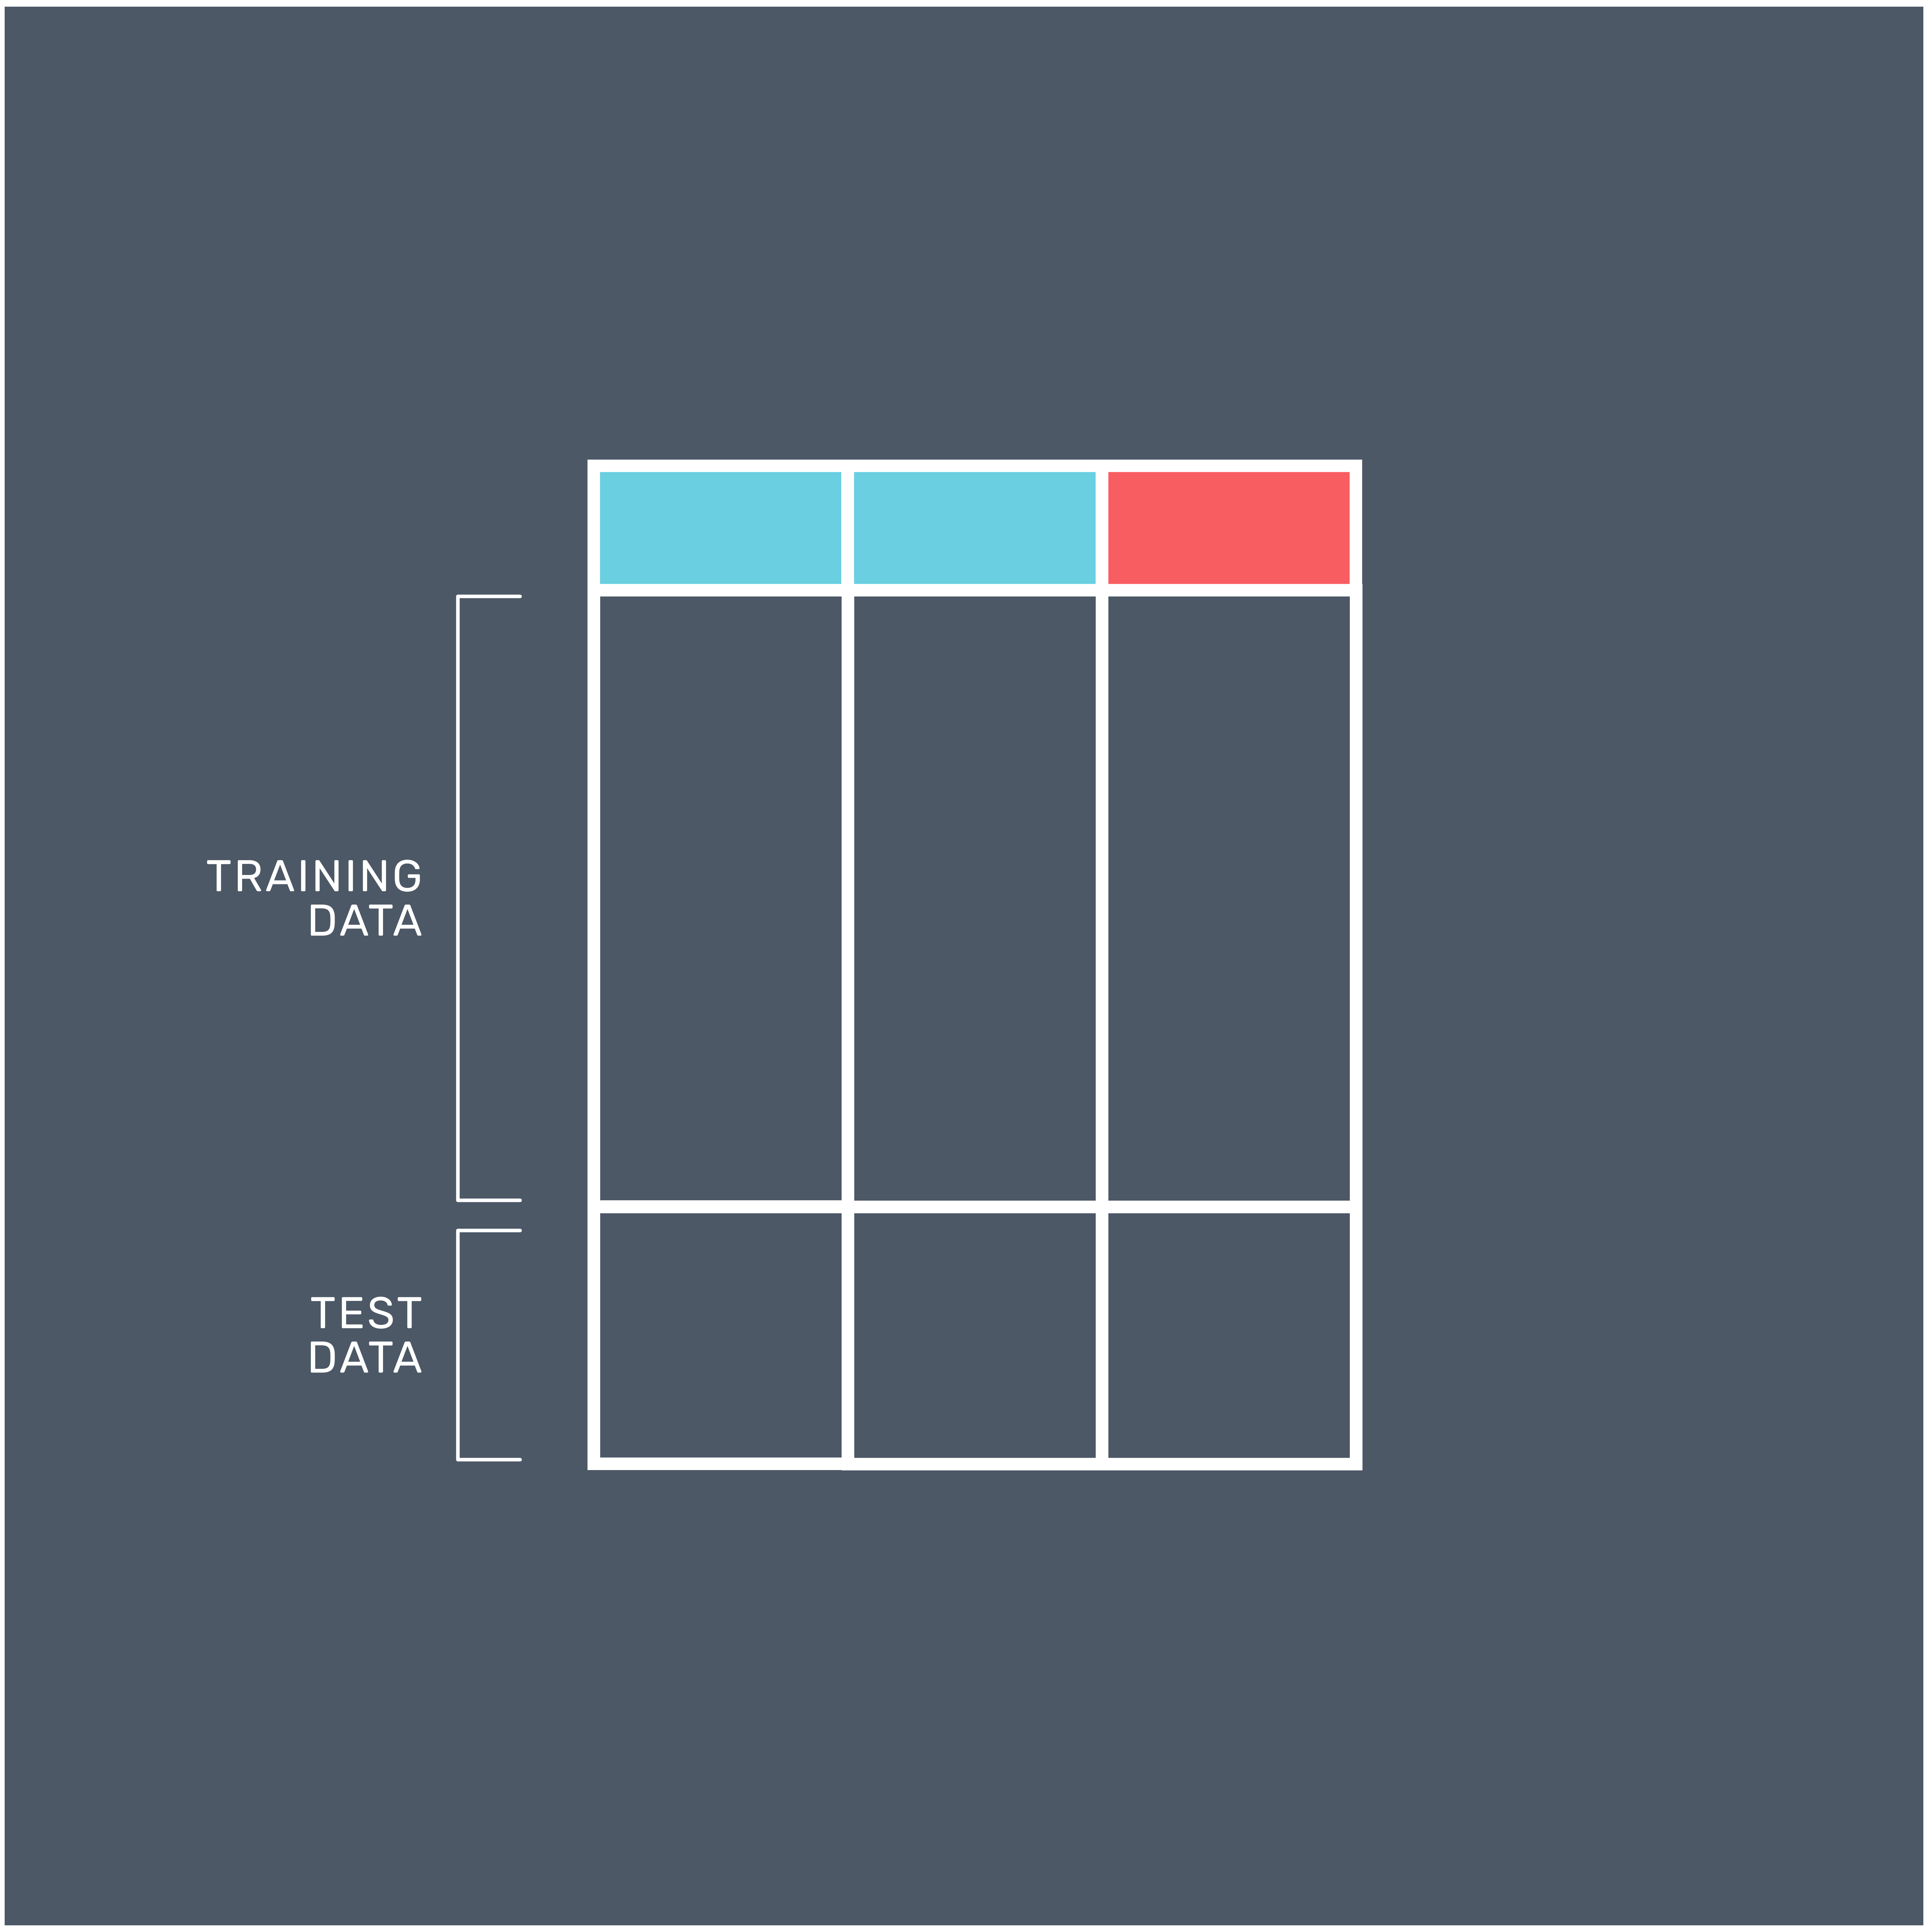

TESTING

Once the training is complete, we need a way to know how the model is performing.

For this, we'll keep a portion of the dataset for *testing*.

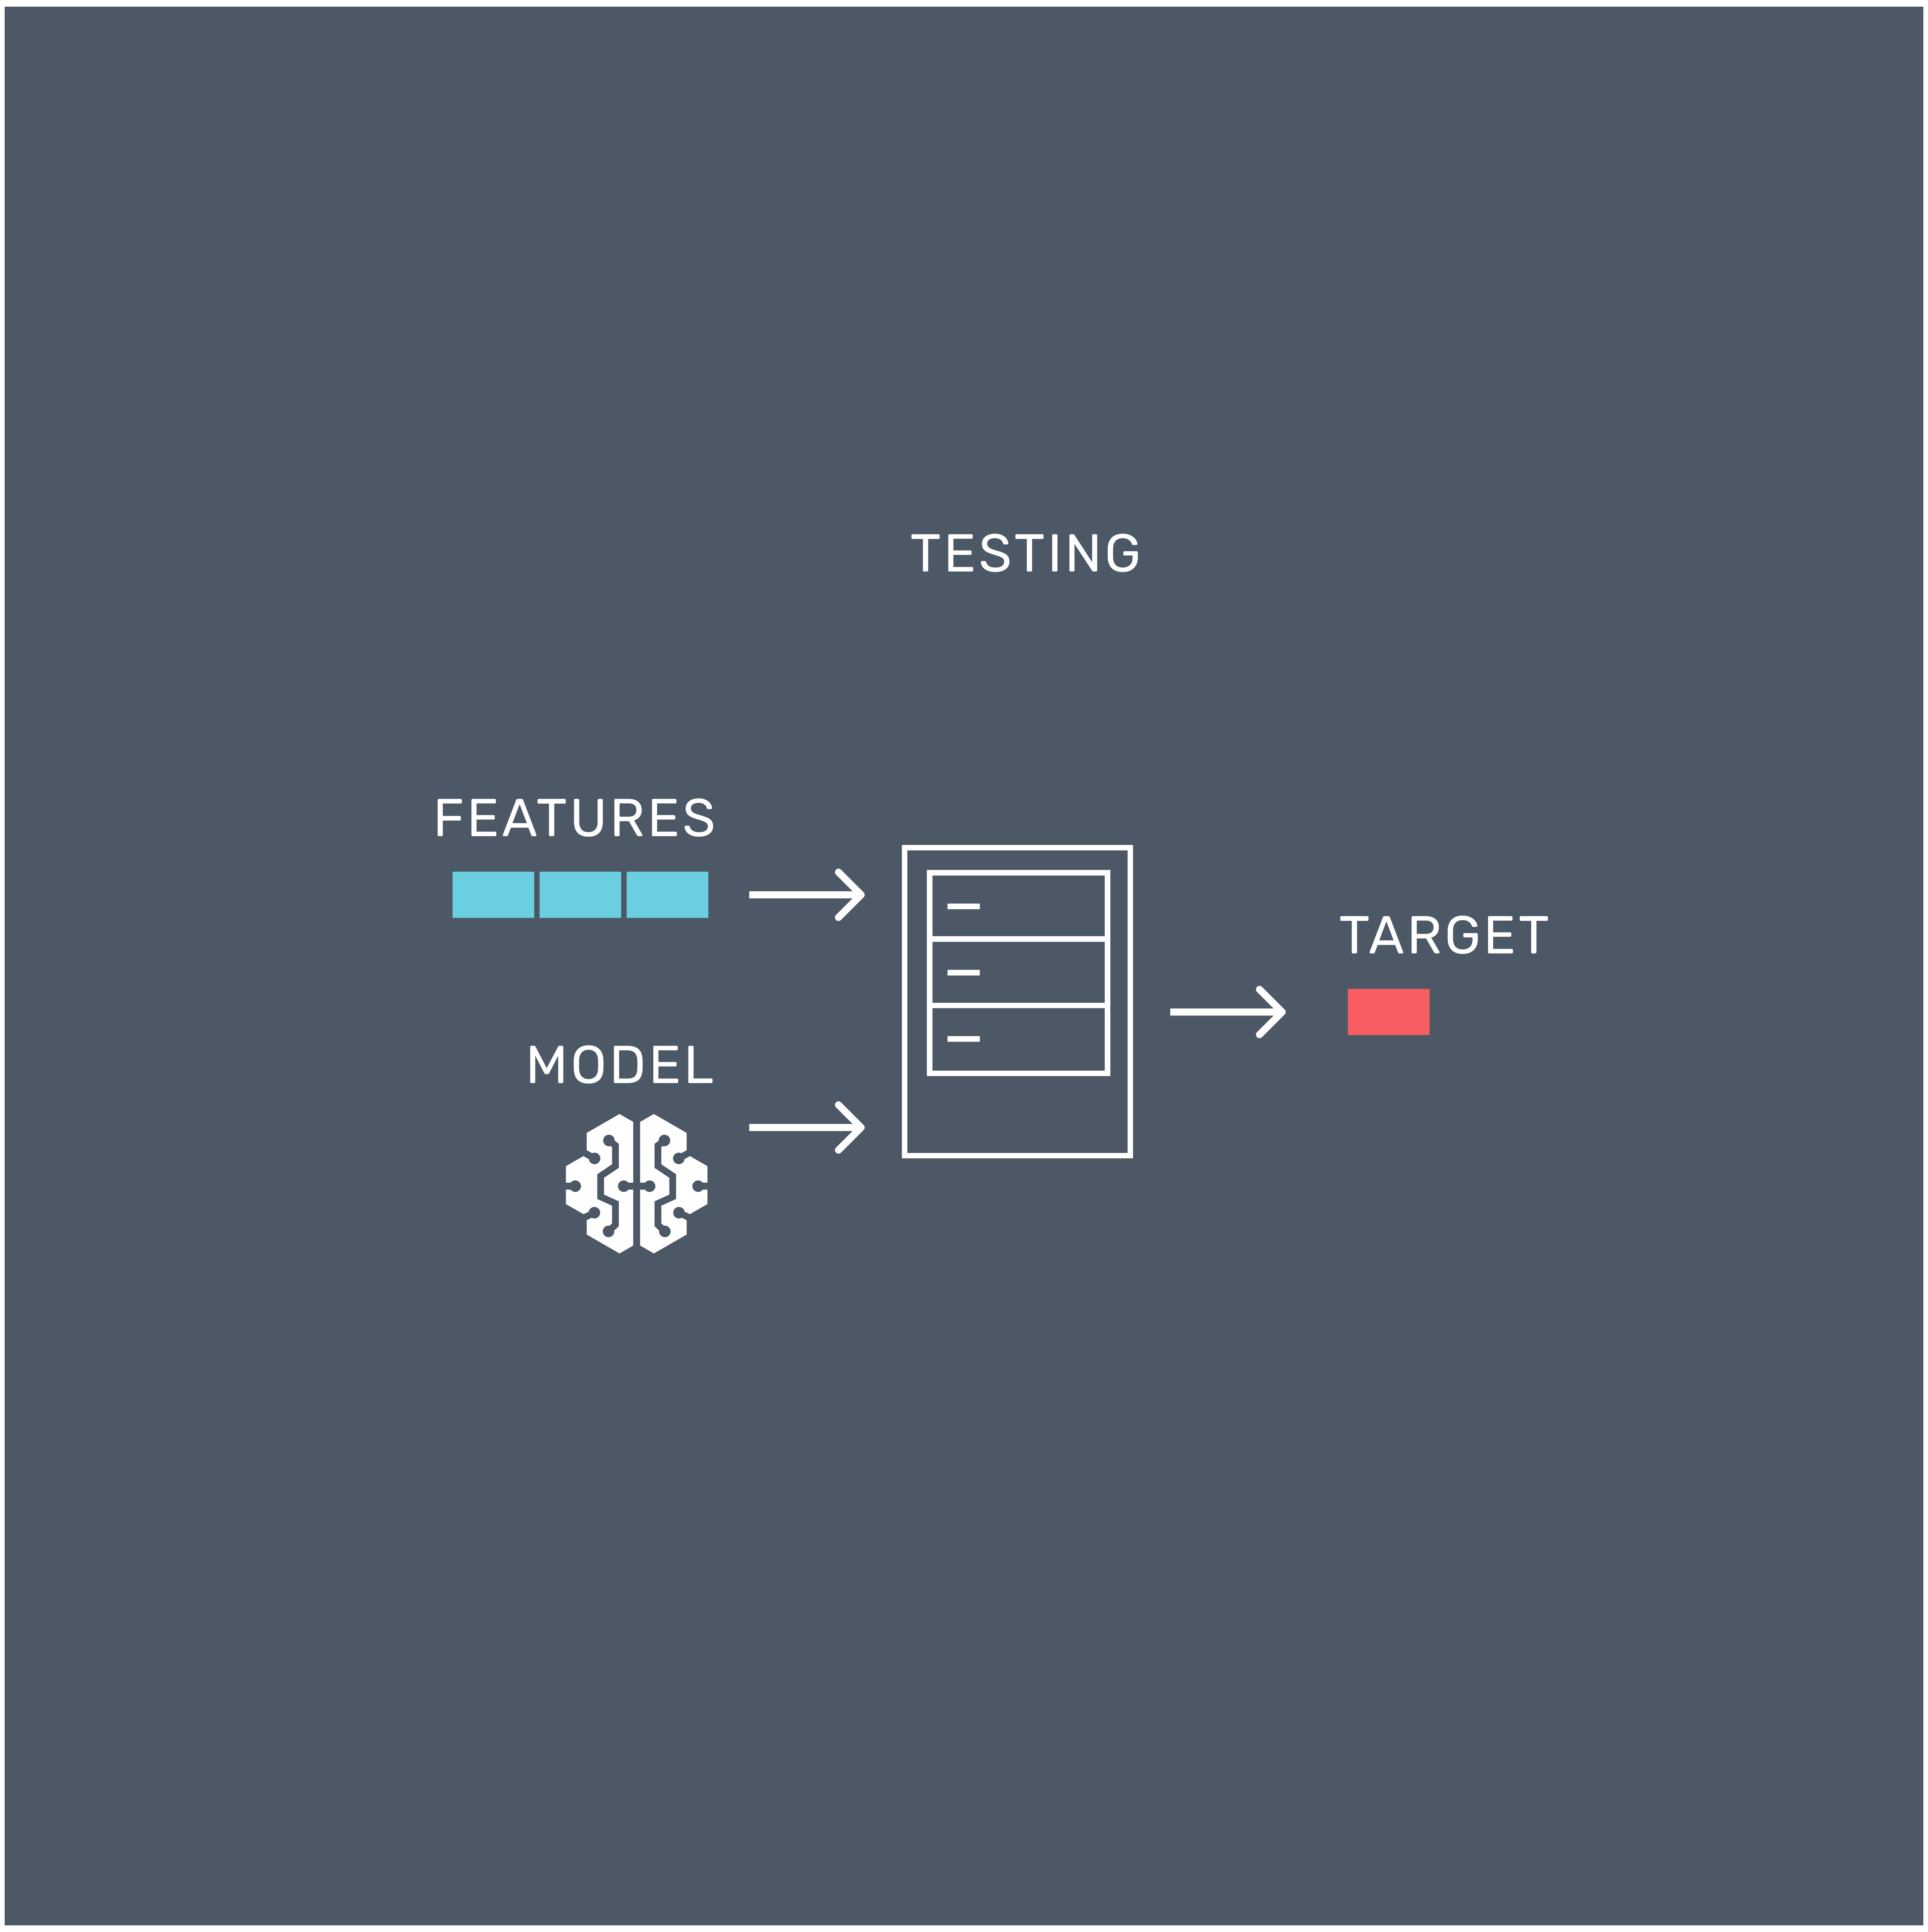

TESTING

During testing, we'll provide the neural network with just the features, without the target. Now that it’s already trained, it’s job is to predict the target values.

In the coming four chapters, we'll revisit these training and testing steps.

GET THE FULL BOOK

Thank you so much for reading the sample of this book. I hope you have found it to be useful.

If you would like to purchase the full copy of the book, you can do so by clicking on the button below.

GET THE FULL BOOK

PRAISE FOR THE BOOK

"*This book takes an impressive no frills approach for people who want to learn about the underpinnings of neural networks in the most time-effective way possible.*"

— **Sebastian Raschka, Ph.D.** (author of Python Machine Learning)

"*This book is like a CEO summary of deep learning. This is an ideal introduction for people who have limited time but still want to go beyond trivial, hand-waving explanations about the core concepts in deep learning.*"

— **Ronald T. Kneusel, Ph.D.** (author of Practical Deep Learning: A Python-Based Introduction and Math for Deep Learning)

CONTACT

I'd love to know if you have any feedback or questions. If you do, you can e-mail them to contact@kdimensions.com.

Thanks again for reading!

Meor Amer

kDimensions