ABOUT THE AUTHOR

Dr Camilla Pang holds a PhD in biochemistry from University College London and is a postdoctoral scientist specializing in translational bioinformatics. At the age of eight, Camilla was diagnosed with Autistic Spectrum Disorder (ASD), and ADHD at twenty-six years old. Her career and studies have been heavily influenced by her diagnosis and she is driven by her passion for understanding humans, our behaviours and how we work.

Instagram @millie_moonface
Twitter @millzymai
www.camillapang.com

Explaining Humans

What Science Can Teach Us about Life, Love and Relationships

DR CAMILLA PANG

PENGUIN BOOKS

Contents

Introduction

It was five years into my life on Earth that I started to think I'd landed in the wrong place. I must have missed the stop.

I felt like a stranger within my own species: someone who understood the words but couldn't speak the language; who shared an appearance with fellow humans but none of the essential characteristics.

In our garden at home I would sit in a multicoloured tent tilted sideways – my spaceship – with an atlas laid out in front of me, wondering what it would take to blast off back to my home planet.

And when that didn't work, I turned to one of the few people who maybe did understand me.

'Mum, is there an instruction manual for humans?'

She looked at me blankly.

'You know . . . a guidebook, something that explains why people behave the way they do?'

I can't be certain – picking up on facial expressions was not, is not and never has been my forte – but in that moment I think I saw my mother's heart break.

'No, Millie.'

It didn't make sense. There were books on almost everything else in the universe, but none that could tell me how to *be*; none that could prepare me for the world; none that could teach me to place a comforting arm around the shoulder of

stereotypically autistic, and too weird to be neurotypically normal, I see myself as an interpreter between both worlds in which I have lived.

I also know that what changed my life was being aware that I was seen and understood. Realizing that I was a person, and had the right to be myself: in fact the duty to be. Everyone has the right to human connection – to be heard and taken seriously. Especially those who, by nature and instinct, struggle to connect. I hope through all the experiences and ideas I share in this book, I will be able both to emphasize the importance of common ground between us as people, and to offer new thoughts on how to achieve it.

So I invite you to join me, on this journey into the strange world of my Aspergic, ADHD brain. It's an odd place to be, but certainly never dull. As well as a notebook, pack your headphones – mine rarely leave my ears, a useful barrier between me and the sensory overload of the outside world. And with that, you're ready. Let's go.

1. How to (actually) think outside the box

Machine learning and decision making

'You can't code people, Millie. That's basically impossible.'

I was eleven, and arguing with my older sister. 'Then how do we all think?'

It was something I knew instinctively then, but would only come to understand properly years later: the way we think as humans is not so different from how a computer program operates. Every one of you reading this is currently processing thoughts. Just like a computer algorithm, we ingest and respond to data – instructions, information and external stimuli. We sort that data, using it to make conscious and unconscious decisions. And we categorize it for later use, like directories within a computer, stored in order of priority. The human mind is an extraordinary processing machine, one whose awesome power is the distinguishing feature of our species.

We are all carrying a supercomputer around in our heads. But despite that, we get tripped up over everyday decisions. (Who hasn't agonized over what outfit to wear, how to phrase an email or what to have for lunch that day?) We say we don't know what to think, or that we are overwhelmed by the information and choices surrounding us.

That shouldn't really be the case when we have a machine as powerful as the brain at our disposal. If we want to improve

how we make decisions, we need to make better use of the organ dedicated to doing just that.

Machines may be a poor substitute for the human brain – lacking its creativity, adaptability and emotional lens – but they can teach us a lot about how to think and make decisions more effectively. By studying the science of machine learning, we can understand the different ways to process information, and fine-tune our approach to decision making.

There are many different things computers can teach us about how to make decisions, which I will explore in this chapter. But there is also a singular, counter-intuitive lesson. To be better decision makers, we don't need to be more organized, structured or focused in how we approach and interpret information. You might expect machine learning to push us in that direction, but in fact the opposite is true. As I will explain, algorithms excel by their ability to be unstructured, to thrive amid complexity and randomness and to respond effectively to changes in circumstance. By contrast, ironically, it is we humans who tend to seek conformity and straightforward patterns in our thinking, hiding away from the complex realities which machines simply approach as another part of the overall data set.

We need some of that clear-sightedness, and a greater willingness to think in more complex ways about things that can never be simple or straightforward. It's time to admit that your computer thinks outside the box more readily than you do. But there's good news too: it can also teach us how to do the same.

Machine learning: the basics

Machine learning is a concept you may have heard of in connection with another two words that get talked about a lot – artificial intelligence (AI). This often gets presented as the next big sci-fi nightmare. But it is merely a drop in the ocean of the most powerful computer known to humanity, the one that sits inside your head. The brain's capacity for conscious thought, intuition and imagination sets it apart from any computer program that has yet been engineered. An algorithm is incredibly powerful in its ability to crunch huge volumes of data and identify the trends and patterns it is programmed to find. But it is also painfully limited.

Machine learning is a branch of AI. As a concept it is simple: you feed large amounts of data into an algorithm, which can learn or detect patterns and then apply these to any new information it encounters. In theory, the more data you input, the better able your algorithm is to understand and interpret equivalent situations it is presented with in the future.

Machine learning is what allows a computer to tell the difference between a cat and a dog, study the nature of diseases or estimate how much energy a household (and indeed the entire National Grid) is going to require in a given period. Not to mention its achievements in outsmarting professional chess and Go players at their own game.

These algorithms are all around us, processing unreal amounts of data to determine everything from what film Netflix will recommend to you next, to when your bank decides you have probably been defrauded, and which emails are destined for your junk folder.

Although they pale into insignificance to the human brain, these more basic computer programs also have something to teach us about how to use our mental computers more effectively. To understand how, let's look at the two most common techniques in machine learning: supervised and unsupervised.

Supervised learning

Supervised machine learning is where you have a specific outcome in mind, and you program the algorithm to achieve it. A bit like some of your maths textbooks, in which you could look up the answer at the back of the book, and the tricky part was working out how to get there. It's supervised because, as the programmer, you know what the answers should be. Your challenge is how to get an algorithm to always reach the right answer from a wide variety of potential inputs.

How, for instance, can you ensure an algorithm in a self-driving car will always recognize the difference between red and green on a traffic light, or what a pedestrian looks like? How do you guarantee that the algorithm you use to help diagnose cancer screens can correctly identify a tumour?

This is classification, one of the main uses of supervised learning, in which you are essentially trying to get the algorithm to correctly label something, and to prove (and over time improve) its reliability for doing this in all sorts of real-world situations. Supervised machine learning produces algorithms that can function with great efficiency, and have all sorts of applications, but at heart they are nothing more than very fast sorting and labelling machines that get better the more you use them.

Unsupervised learning

By contrast, unsupervised learning doesn't start out with any notion of what the outcome should be. There is no right answer that the algorithm is instructed to pursue. Instead, it is programmed to approach the data and identify its inherent patterns. For instance, if you had particular data on a set of voters or customers, and wanted to understand their motivations, you might use unsupervised machine learning to detect and demonstrate trends that help to explain behaviour. Do people of a certain age shop at a certain time in a certain place? What unites people in this area who voted for that political party?

In my own work, which explores the cellular structure of the immune system, I use unsupervised machine learning to identify patterns in the cell populations. I'm looking for patterns but don't know what or where they are, hence the unsupervised approach.

This is clustering, in which you group together data based on common features and themes, without seeking to classify them as A, B or C in a preconceived way. It's useful when you know what broad areas you want to explore, but don't know how to get there, or even where to look within the mass of available data. It's also for situations when you want to let the data speak for itself, rather than imposing pre-set conclusions.

Making decisions: boxes and trees

When it comes to making decisions, we have a similar choice to the one just outlined. We can set an arbitrary number of possible outcomes and choose between them, approaching

problems from the top down and starting with the desired answer, much like a supervised algorithm: for example, a business judging a job candidate on whether they have certain qualifications and a minimum level of experience. Or we can start from the bottom, working our way upwards through the evidence, navigating through the detail and letting the conclusions emerge organically: the unsupervised approach. Using our recruitment example, this would see an employer consider everyone on their merits, looking at all the available evidence – someone's personality, transferable skills, enthusiasm for the job, interest and commitment – rather than making a decision based on some narrow, pre-arranged criteria. This bottom-up approach is the first port of call for people on the autistic spectrum, since we thrive on bringing together precisely curated details to form conclusions – in fact we need to do that, going through all the information and options, before we can even get close to a conclusion.

I like to think of these approaches as akin to either building a box (supervised decision making) or growing a tree (unsupervised decision making).

Thinking in boxes

Boxes are the reassuring option. They corral the available evidence and alternatives into a neat shape where you can see all sides, and the choices are obvious. You can build boxes, stack them and stand on them. They are congruent, consistent and logical. This is a neat and tidy way to think: you know what your choices are.

By contrast, trees grow organically and in some cases out of control. They have many branches and hanging from those

are clusters of leaves that themselves contain all sorts of hidden complexity. A tree can take us off in all sorts of directions, many of which may prove to be decisional dead ends or complete labyrinths.

So which is better? The box or the tree? The truth is that you need both, but the reality is that most people are stuck in boxes, and never even get onto the first branch of a decision tree.

That certainly used to be the case with me. I was a box thinker, through and through. Faced with so many things I didn't and couldn't understand, I clung to every last scrap of information I could get my hands on. In between the smell of burnt toast on weekdays at 10.48 a.m. and the sound of schoolgirls gossiping in cliques, I would engage within my recreational equivalent – computer gaming and reading science books.

Night after night, throughout the years of boarding school, I would revel in my solitude by reading and copying selective bits of texts from science and maths books. My trusty instruction manuals. I took great pleasure and relief from doing this over and over, with different science books, not knowing why but only to reach the crescendo of pinning down some gravitational understanding of the reality before me. My controllable logic. The things I read helped give me rules that I set in stone, from the 'right' way of eating to the 'right' way to talk to people and the 'right' way to move between classrooms. I got stuck in a rut of knowing what I liked and liking what I knew – regurgitating a series of 'should's to myself because they felt safe and reliable.

And when I wasn't sitting with my books, I was observing: memorizing number plates on car journeys, or sitting around dinner tables contemplating the shape of people's fingernails.

As an outsider at school, I would regularly use what I now understand to be classification to understand new people entering my world. Where were they going to fit into this world of unspoken social rules and behaviours that I struggled to understand? What group would they gravitate towards? Which box could I put them in? As a young child I even insisted on sleeping in a cardboard box, day and night, enjoying the feeling of being cocooned in its safe enclosure (with my mum passing biscuits to me through a 'cat flap' cut in the side).

As a box thinker I wanted to know everything about the world and people around me, comforting myself that the more data I accumulated, the better decisions I would be able to make. But because I had no effective mechanism for processing this information, it simply piled up in more and more boxes of useless stuff: like the junk that hoarders can't bear to throw out. I would become almost immobilized by this process, at times struggling to get out of bed because I was so focused on what exact angle I should hold my body at. The more boxes of irrelevant information piled up in my mind,

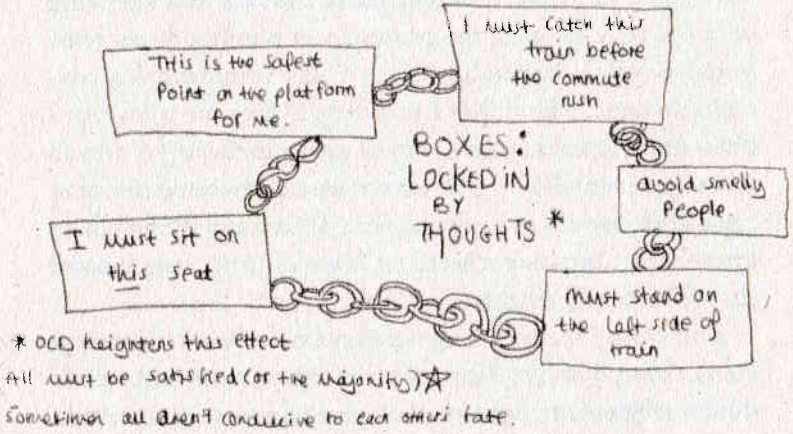

the more directionless and exhausted I became, as every box in my mind started to look the same.

My mind would also interpret information and instructions in a wholly literal way. One time I was helping my mum in the kitchen, and she asked me to go out and buy some ingredients. 'Can you get five apples, and if they have eggs get a dozen.' You can imagine her exasperation when I returned with twelve apples (the shop had indeed stocked eggs). As a box thinker, I was incapable of escaping the wholly literal bounds of an instruction like that, something I still struggle with today: such as my belief, until recently, that one could actually enrol at the University of Life.

Classification is a powerful tool, and useful for making immediate decisions about things, such as which outfit to wear or what film to watch, but it places severe limitations on our ability to process and interpret information, and make more complex decisions by using evidence from the past to inform our future.

By trying to classify our lives, thinking in boxes, we close off too many avenues and limit the range of possible outcomes. We know only one route to work, how to cook just a few meals, the same handful of places to go. Box thinking limits our horizons to the things we already know, and the 'data' in life we have already collected. It doesn't leave much space for looking at things differently, unshackling ourselves from preconceptions, or trying something new and unfamiliar. It's the mental equivalent of doing exactly the same thing at the gym every session: over time your body adapts and you see less impressive results from your workout. To hit goals, you have to keep challenging yourself and get out of the boxes that close in on you the longer you stay in them.