Chapter 1

Introduction

There are two key characters in the story of control, namely, the master puppeteer and the puppet. Assume now that the master puppeteer is us, and our goal is to control the behaviour of the puppet. The puppet can represent any system that requires our control. For instance, the puppet could represent a system that is *biological*, like keeping our physical fitness levels up so that we can run a marathon. It could be *economic*, like a stock market in which we are maximizing our profit by buying and selling shares. It could be *organizational*, like marshalling a troop of soldiers to protect a safety zone. It could even be *ecological*, for instance trying to sustain an endangered ecosystem like a coral reef. More typically, when we think of control systems, what comes to mind is something *industrial* like operating a nuclear power plant, or *mechanical* like driving a car, or *safety critical* like flying a plane.

Clearly, then, there are many examples of control systems, some of which we are likely to experience on a regular basis in our daily lives. Control systems, therefore, are a rather broad subject and, as

Controlling Uncertainty: Decision Making and Learning in Complex Worlds, Magda Osman.

will be made clear in this book, can encompass almost everything. So how are we able to do it; how are we able to exert control over these various types of systems? For those who have thought of this question, this book will be a guide to some answers.

Puppets, Puppets Everywhere …

The reason that control systems can be seen everywhere is that almost anything can be thought of as a system, and most things we do in our lives are driven by our need to control. But, while the systems may be incredibly pervasive, they do adhere to some basic characteristics. Control systems are complicated. This is because they have a number of elements that will vary all at once and from one point in time to another – like puppets. Puppets can have multiple components; their moving parts (e.g., arms, legs, feet, hands fingers and head) are more often than not interconnected (e.g., fingers to hand, hand to arm etc.), and can vary all at once (e.g., performing a jump) as well as singularly in one point in time (e.g., raising a hand to wave). Thus, given these capacious characteristics of control systems, they are quite literally everywhere.

… It Takes All the Running You Can Do, to Keep in the Same Place[1]

In these systems we are often required to manage the events that occur in a way that leads to something predictable, and desirable. This can be incredibly difficult to achieve and takes many years of training (e.g., becoming a pilot of a passenger jet), because the system varies of its own accord, as well as because acting upon it makes it change in some way. Or, more often than not, it is a combination of us acting on it and it doing something itself that produces changes in events. To bring the analogy of the puppet to

[1] Lewis Carroll, *Through the looking glass* (1871).

bear more obviously with control systems, in the story the puppet had its own internal mechanism that also made it move. Imagine how hard it is to control a malleable puppet and make it dance in time to a tune without an internal mechanism that can make it move on its own. Now imagine how much harder it is when it can move on its own and not always predictably. As hard as it seems, we are capable of achieving this. So we return to the question again: how are we able to exert control over such a complicated situation?

Vladimir: 'Say Something!' Estragon: 'I'm Trying. … In the Meantime Nothing Happens'. Pozzo: 'You Find It Tedious?' Estragon: 'Somewhat'.[2]

To answer the question 'How?', we need to find a better way of asking it. First of all, finding some way of describing how these different types of systems work is of great importance, particularly if they can, on a basic level, be thought of in a similar way. Second, to complement this, our ability to control what happens in these systems should reduce to some basic psychological learning and decision-making mechanisms. They should do this because we need psychological mechanisms in place that enable us to predict the behaviour of the system and coordinate our own behaviours to effect a specific change in it. Therefore, finding some way of describing our psychological processes, along with describing the control system itself, is crucial to having an understanding of control (i.e., the scientific pursuit) and being able to improve our ability to manipulate our environment (i.e., the applied pursuit).

Given the extensiveness of both objectives, typically at the start of books like this there is a tendency to spell out at the beginning what things will not be included and what can't be achieved. I am going to avoid this. The aim of this book is to be as inclusive as possible. If you've flicked through it already, you will have

[2] Samuel Beckett, *Waiting for Godot* (1954/2009).

noticed that there are chapters spanning subject areas that include philosophy, engineering, cybernetics, human factors, social psychology, cognitive psychology and neuroscience. In order to get to the answer of 'How?', we need to consider the various contributions that each of these subjects has made. The issue of control invites attention from many disciplines that don't always speak to each other. Putting them side by side in chapters in a book is also a way of showing how they in fact do relate. Moreover, they also provide the groundwork for my answer to the question of 'How?' which is presented at the end of this book.

There are two important ideas that will help to carry you along this book: (1) all the themes introduced in this book are reducible to five basic concepts: control, prediction, cause–effect associations, uncertainty and agency; and (2) all of the issues that these basic concepts raise are ultimately, and will in this book be, directed towards addressing one question, which for the purposes of this book is THE question: *how do we learn about, and control online, an uncertain environment that may be changing as a consequence of our actions, or autonomously, or both?*

To understand the issue of control psychologically, and to understand the control system itself in all its various guises, we have to become familiar with these five core concepts and how they are tackled through the eyes of each of the aforementioned subjects. However, I am not alone; this endeavour has been embarked on by many,[3] and throughout the different chapters of the book it will become apparent that there are various ways of understanding the psychological and objective characteristics of control systems. Therefore, I will take the opportunity here to qualify why this book is not a reinvention of the wheel, by stating what it hopes to do differently.

[3] Cybernetics is the best-known example of an interdisciplinary movement designed to examine all issues related to control and self-regulation (see Chapter 4). More recently, machine-learning theorists have also attempted to draw work from engineering, biology, psychology and neuroscience to develop formal descriptions of behaviours associated with learning and controlling outcomes (e.g., Sutton & Barto, 1998).

The Aim of This Book

Role 1: catalogue

At its most humble, this book serves the purpose of being an inventory of sorts of what we currently know in a range of disciplines (e.g., engineering, artificial intelligence [AI], human factors, psychology and neuroscience) about control systems and control behaviour. Though not ever seriously taken up, an appeal of this kind was made in the late 1940s by Wiener, the self-proclaimed father of cybernetics – a discipline designed to study all matters related to self-organizing systems. Wiener (1948) hoped to bring together many disciplines to understand common problems concerning control. However, Wiener (1948) proposed that 'the very speed of operations of modern digital machines stands in the way of our ability to perceive and think through the indicators of danger' (p. 178). The effort in understanding all matters related to control came with a warning that technological advances may be such that artificial autonomous agents would be controlling our lives. That is, in the future the puppet would eventually rule the puppeteer, and not the other way around. Though the worry that control systems will reach a level of self-organization that may challenge our mastery of the world is perhaps unwarranted, surveying the most recent advances in theory and practice should give us a better understanding of what control systems can do, and our place with respect to them.

As suggested, an overhaul of this kind has yet to be undertaken, and so this book is an opportunity to do just that. For instance, due to the increasing complexity of the systems under our control (e.g., systems that identify tumours in X-ray images, voice recognition, predicting stock market trends, creating game play in computer games and profiling offenders), there are in turn ever increasing demands placed on them to achieve optimal performance reliably. Even something as prosaic as the car now includes an increased level of automation. This is generically classified under the title of *driver assist systems* (DAS). DAS now include electric power-assisted steering (EPAS), semi-automatic parking (SAP), adaptive cruise control (ACC), lane departure warning (LDW) and vehicle stability control (VSC). All of these things now influence the ride and handling of vehicles we drive. So we might ask ourselves, if we have handed over so much autonomy to the car, what control do we have?

More to the point, disciplines such as control systems engineering present us with ever growing challenges because the control systems (e.g., car) that are part of our everyday interactions continue to increase in their capabilities and complexity. If complexity is increasing, then surely we need to know how we cope with it now, especially when things go wrong. Increasing complexity in our everyday lives doesn't just come from controlling devices such as cars. There has been a charted increase in the complexity of the decision making involved in economic, management and organizational domains (Willmott & Nelson, 2003). We can spot this complexity because some of it filters down to our consumer choices. For instance, take shopping. We have to adapt to the growing complexity that we face in terms of the information we have to process (e.g., more available product information), the choices we are presented with (e.g., more products to choose from) and the changing goals that we are influenced by (i.e., desires, aspirations and expectations). At the heart of adapting to the increasing level of complexity in our lives is our ability to still exert control. So, given the new challenges and demands that are placed on us in our lives right now, this book may be considered a sort of stock take of relevant and current research in the study of all things control related.

Role 2: solving the problem of complexity

A broader aim of the book is to help clarify what we mean when we say an environment is complex, and what it is about control systems that invites researchers from different disciplines to refer to them as complex. The complexity issue is important for the reason that there needs to be a cohesive idea about what makes control systems difficult to understand, and why we can fall into traps when we come to control them.

For instance, the term *complexity* has a specific reference in design engineering and is a measure of the structure, intricateness or behaviour of a system that characterizes the relationship between its various components.[4] Thus, the properties of a control system can be specified according to objective characteristics of complexity from an engineering perspective. Efforts in defining complexity have also been attempted from a psychological perspective. Studies of human behaviour in control systems have taken properties of a control system (e.g., transparency, dynamics, number of variables, number of connections and functional forms – linear, curvilinear, stochastic and feedback loops) and investigated how competent we are at controlling systems when these properties are manipulated. The steady amassing of data from experiments along these lines started from early work by Dörner (1975). But, unfortunately, despite the wealth of findings, there has been little headway in being able to say generally what contributes to making a system complex from a psychological point of view.

This is a major problem. We simply can't say what makes systems complex *in general*. What we can only say is what might make a *particular* system complex. This is hugely limiting because we can't generalize to different types of control systems the psychological factors associated with them. In other words, the analogous situation would be this: we might know that the puppet in your hand is going to be hard work to operate because every joint is movable. But we wouldn't know why it is that if, given a variety of other puppets, you still find it hard to operate them all. Is it down to psychology? Is it down to the way the different puppets function? Or is it a combination of both?

Attempts to answer these questions involve identifying possible measures of control systems complexity in order to predict the

[4] As systems become more sophisticated in their functions, the interrelationship of the controlled variables in turn is used as a measure of the systems' complexity.

success of psychological behaviours. Again, these include borrowing ideas from another discipline, in this case computer science, which have been used to measure how controllable the system is (e.g., whether it is in non-polynomial time [NP] or Polynomial time [P]),[5] the size of the system (e.g., its search space) and the number of interdependent processes that are contained within it. However, to date, none of these captures the true complexity of a control system, or for that matter accurately predicts what psychological behaviours we are likely to use. While it might be the case that from a computer science perspective a complicated control system can be described in some mathematical way, and from this, based on a long arduous formal process, some claims are made about how we ought to behave in order to control it, there is a problem. There is a lack of correspondence. Humans have elegant and simple means (e.g., heuristics)[6] to reduce complexity that may be formally difficult to describe, but equally there are examples in which humans find it virtually impossible to tame a complex situation that may, from a computer science perspective, be mathematically simple to describe.

Thus, there is a gap. Knowing what defines the system as complex is no guarantee for understanding the types of behaviours needed to learn about it in order to eventually control it. Therefore, the aims of the book are to present ways of clarifying the problems faced when attempting to define complexity, and to offer a solution to it. The solution is based on describing control systems as uncertain environments. What I propose is that identifying and measuring uncertainty of the system boil down to tracking over time when changes to events in the system occur that are judged to occur independently of our actions, while also tracing those changes that are rare but have a substantial impact on what happens overall, and

[5] This relates to a general issue concerning the solvability and verification of a problem and the number of computational steps that are needed (i.e., polynomial time) to achieve both.

[6] A method by which prior knowledge and experience are used to help quickly generate decisions and actions to solve a problem.

that also didn't result from our actions. These descriptions of the system's behaviour need to be integrated with people's judgements about how confident they are that they can predict the changes that occur in a control system, and people's judgements as to how confident they are that they can control the changes that occur in a control system. These are what I claim to be the bare essentials of understanding and developing a metric for examining the success of controlling uncertainty in control systems. However, you could argue that I've just performed a sleight of hand by saying complexity is uncertainty in which case I haven't solved the problem of complexity at all. Let this might be true, but let me qualify a few things. First of all, I'm not claiming that complexity and uncertainty are the same thing. What I am claiming is that translating the issue of complexity into defining what makes a control system uncertain paves the way for connecting objective descriptions of the control system with psychological descriptions of our behaviour. Thus, uncertainty as a concept is better than complexity as a bridge between the control system and us. Second, I'm not saying complexity is reducible to the other. Instead, I am suggesting that, to answer the question of what makes a system complex a new perspective ought to be taken. From a practical standpoint, attempts to solve the mystery of what makes a system complex have thus far been illuminating but ultimately unsuccessfully in defining what makes a system difficult for us to control. I argue that it may be easier to decide on the objective properties of the control system that are uncertain than those that make it complex. This pragmatic point then, arguably, is a good starting point for presenting a framework that develops on the solution to the problem of complexity. Better still would be a framework that also tries to bridge objective uncertainty with psychological uncertainty. Here, then, is the final ambitious objective of this book.

Role 3: providing answers to how we control uncertainty

In an effort to capitalize on the humble and broad objectives of this book, the final aim is to present some general principles along with a framework that relates work focused on designing specific kinds of control systems (such as engineering and machine learning), work examining the interaction between humans and control systems (human factors, human–computer interaction [HCI] and ergonomics), and work investigating the underlying psychological mechanisms that support behaviours that enable control (social psychology, cognitive psychology and neuropsychology). By taking current stock and by tackling a number of critical issues that have dogged many fields, there are new insights to be gained. At the very least, in pursuing this objective a number of new questions can be raised and ways of answering them can now be formulated. This discussion is reserved for the final chapter, and can also be viewed as a sort of roundup of the main issues surrounding the puppet and the puppeteer. Moreover, all the details needed to understand the framework are crystallized by the puppeteer's final comment to the ringmaster.

> 'All I know is this', said the puppeteer. 'To make the puppet dance the way I wanted, I had to know how supple the limbs and head were. The suppler it seemed to be, the more elegantly I could make it dance. I had to know when the mechanism in the puppet was running fast or slow, and when it failed to behave as I'd expected from what I did. This helped me to decide the times when it simply seemed that it was indeed behaving of its own accord, from times that it was behaving as I wanted. For all this, I needed to keep in mind two very important details: that I was the one who could choose what levers to operate, and I was the one who could choose when to operate them. This is how I gathered what the puppet might do from one moment to the next, and this was what helped me to understand what more I needed to know to make it dance to different tunes'.

Whether or not the final objective is successfully achieved is left to be judged by the reader. But one can read all the related and current work on all matters concerning control systems and control behaviours in Chapters 2–8 without relying on the proposed framework to do so.

The Structure of This Book

To provide as comprehensive an account as possible of the work that relates to control systems and control systems behaviour, a broad range of disciplines are covered by this book. In short, if we go back to the story of control to understand the structure of the book, to begin we take the point of view of the ringmaster who sees everything: the puppet in its box, the strings, the panel with the levers and the puppeteer. From this point of view, we can consider all the important issues in the situation. Next we focus on the puppet and its internal mechanisms, including how it operates and what it can do. From there we consider the strings that link the puppet to the puppeteer. Finally, we need to take into account the puppeteer, and the internal mechanisms and processes that make up his behaviour. By following this order we can build up a picture of the whole story of control from puppet to puppeteer, and from this we can see how we get to the answer at the end.

So, to begin, the fundamental issues that concern control are grounded in philosophy, namely, the way in which we construe causality and our sense of agency. Both of these are central to understanding how we interact with, and assert control over, our environment. The next step is to consider the control systems environment itself in terms of its more common reference to engineering. In particular, we will examine how engineers develop control systems, and what methods they use to formally describe them. This helps to illuminate the general aspects of the environment worth knowing about which apply to all other types of control systems. Closely related to engineering are the fields of cybernetics, AI and machine learning. To situate these fields of research in the context of the general issues they tackle, and the historical basis for them, the next chapter spends some time discussing cybernetics. Cybernetics, as mentioned before, was designed to provide a framework for understanding issues of control in all possible self-regulatory environments, and in so doing, it also focused on examining how humans interact with and relate to them. This has become a research matter that has gained in importance, and human factors research specifically tackles this. Thus the chapter on this subject serves as the intersection between the first and second halves of this book. The second half of the book is orientated towards control behaviours from a psychological perspective. Across the next three chapters, the discussion focuses on general descriptions of our behaviour at a social level, and the underlying mechanisms that support that behaviour at a cognitive and neuropsychological level. The final chapter draws on all the work that has been presented throughout the book in order to end with some general principles of the control system and the psychological processes involved in it. This is the basis for presenting a succinct framework that takes into account the properties of the environment and the ways in which we as puppeteers master it.

To help the reader, below is a short summary of each chapter, so that the general gist of each is conveyed and can more easily aid the reader in deciding what may be relevant to refer to.

Chapter 2: Causation and agency

This chapter considers our sense of agency (i.e., the relationship between thoughts, beliefs, intentions, desires [reasons] and the events that occur in the external world [actions]) and causality, our experiences and understanding of relations between events. Much of the knowledge we develop while interacting with our environment, particularly one which is uncertain, is anchored by the actions we generate, and the sense of agency attached to them. This helps us to relate what we do with the effects that are produced. Thus, understanding the deeper issues related to these basic assumptions provides an important foundation for tackling the issues that will be discussed throughout the book.

Chapter 3: Control systems engineering

Engineering work on control systems has a long tradition dating back to the nineteenth century. Control systems engineering is concerned with how dynamic systems operate, and provides ways of formally describing them. Engineering not only aims to accurately describe the system, but also has another objective, which is to offer a greater level of automaticity in the systems that surround our everyday lives. In order to present the details of the control system from the viewpoint of engineering, the chapter introduces control theory, and discusses how it is applied in order to help develop systems that offer a high level of automation.

Chapter 4: Cybernetics, artificial intelligence and machine learning

Cybernetics has been pivotal in suggesting that what underlies control systems in engineering applies to any systems we can think of in nature and society, whereas the way in which complex devices function and help to offer adaptive control has been the focus of study in engineering. Importantly, the issues faced by engineering are also encountered by machine learning and AI. So, the chapter aims to draw attention to the critical issues that are faced by current research in machine learning and AI, and, in so doing, to highlight the progress made thus far, and the limitations that are currently faced.

Chapter 5: Human factors (HCI, ergonomics and cognitive engineering)

Some of the key concerns for human factors research is what happens when errors in control systems occur, and how things can be improved given the increasing exposure that humans have to automated control systems. What this chapter illuminates is that there are often differences in the kinds of assumptions that designers make about human capabilities, and the assumptions that human operators of control systems make about how the control system behaves. Misalignment of assumptions can be the basis for many of the problems that arise when humans interact with complex systems.

Chapter 6: Social psychology, organizational psychology and management

What this domain of research helps to draw attention to is that motivation, and the ways in which we pursue goals, can have a profound impact on our ability to successfully control our immediate environment. Moreover, agency takes centre stage in research and theory in the study of control in social psychology. This has also been central to understanding psychological behaviours in organizational and management contexts, which are also discussed. Thus the chapter covers the main research findings and current theoretical positions on control in social psychology with examples of how it applies in management and organizational contexts as examples of control systems.

Chapter 7: Cognitive psychology

Cognitive psychology has amassed over 40 years of research examining the underlying mechanisms associated with controlling complex systems. More specifically, work on control systems behaviour is concerned with how we acquire knowledge about the system, and how we apply it to reach different goals. The aim of this chapter is to introduce the types of psychological tasks that have been used to study control behaviours, along with the general findings and theories to account for them. In so doing, the chapter also discusses relevant work on perceptual-motor control, causal learning and reasoning, and predictive judgements in multiple cue probabilistic learning tasks.

Chapter 8: Neuroscience

The chapter approaches neuroscience in two ways: it suggests that neurological functioning of the brain can be thought of as a control system with respect to feedback, dynamics, uncertainty and control. It also considers the neuropsychological basis for behaviours associated with control. By looking at the underlying mechanisms that support decision-making and learning processes, the chapter examines current work that has contributed to understanding uncertainty and the role of feedback from a neuroscientific perspective. To bring all of this to the fore, the latter half of the chapter introduces a new field of research: neuroeconomics. Neuroeconomics has been used as a way to draw together neuroscience, economics and psychology under the common vision of understanding how we learn about and make decisions in uncertain environments.

Chapter 9: Synthesis

The aim of this chapter is to propose that the control system is an uncertain environment, and therefore it is crucial to know whether the effects that occur in it are the result of our own actions or some aspect of the way the system behaves independently of us. This kind of uncertainty can reflect a genuine feature of the environment: that it is highly probabilistic in nature and hard to predict. But uncertainty is not only a feature of the environment; it can also be generated by the individual acting on the system. We can be uncertain about how accurately we can predict the system will behave, and unsure about our ability to effect a change in the system. The chapter lays out a number of general principles from which a framework is described that provides a general account of control.

Chapter 10: Epilogue

The aim here is to come full circle. The story is presented again, but this time I've slotted into various places the five core concepts discussed in this book along with the framework proposed in the final chapter. I hope that this will at least provide a final entertaining glimpse at ways of understanding how we control uncertainty.