

Chapter 1: Introduction

Some scientists invest an entire career in the study of organisms of a single species, others in understanding particular types of cells or in determining the role of a certain gene. The elements of each level of biological organization can take more than a lifetime to understand. How then can we put all this information together? Understand how genes interact to drive the cell, how cells interact to form organisms, and how organisms interact to form groups and societies? These questions are fundamental to the scientific endeavor: how do we use our understanding of one level of organization to inform us about the level above?

Linking different levels of organization involves the study of collective phenomena: phenomena in which repeated interactions between many individuals produce patterns on a scale larger than themselves. Collective phenomena are within us and all around us: the clustering of cells to build our bodies; the firing of neurons in our brains; flocks of birds twisting above our heads; and the pulsating mass of bodies surrounding us on a Saturday night dancefloor. Understanding these phenomena is an important part of the fields of developmental biology, neuroscience, behavioural ecology and sociology, to name just a few. Even researchers studying the most intricate details of the components of a particular system are acutely aware of the need to understand how these components fit together to create a whole system.

The study of collective phenomena is founded on the idea that a set of techniques can be applied to understand systems at many different physical scales. This idea originated from mathematics, theoretical physics and chemistry. Books by Wiener (1948), Ashby (1947), von Bertalanffy (1968) and Nicolis & Prigogine (1977) all aimed at providing a framework for the study of collective phenomena. Von Bertalanffy argued for the existence of general growth laws of social entities as diverse as manufacturing companies, urbanisation and Napoleon's empire. Wiener argued that homeostasis, a stable functioning of natural systems, could be achieved through simple feedback loops. Nicolis and Prigogine aimed to pin down a rigorous theory of non-linear thermodynamics, explaining similarities between systems at very different scales. For example, could the flow of traffic be described by the mathematics of fluid flow? And if this was the case could we make general statements about the flow of any type

of matter, be it swarms of locusts, crowds leaving football grounds or water running down the drain?

Over the last 30 years, research into collective phenomena and complex systems has rapidly expanded. The completion of the human and other genome projects was followed by a call for systems biology, a combination of experimental and computational approaches to integrate our collected database of biological facts (Kitano 2002). The challenges of assessing changes in the global environment require understanding of ecological interactions that occur over many different temporal and spatial scales (Levin 1992; Levin 2000; Stainforth et al. 2005). Study of neurobiology and the immune system again involve understanding how neurons or cells interact to make decisions (Bays & Wolpert 2007; George et al. 2005). Modern sociology and social psychology aim to link the decisions of individuals to the social behaviour of the many (Hedstrom 2005; Milgram 1992; Schelling 1978). Accompanying the realization of the importance of collective phenomena within different fields has been the development of mathematical modelling tools for investigating these systems.

Animal groups provide many key examples of collective phenomena. They also provide some of the most spectacular and fascinating sights in the natural world (see the pictures on the website). Flocks of birds turning in unison or migrating in well ordered formation; fish shoals splitting and reforming as they outmaneuver a predator; swarms of honey bees settled on a branch of a tree while colony members use dances to debate to where they will fly; and the long bifurcating trails along which ants transport food and materials have all long fascinated scientists. There is something captivating about the patterns these groups create. They are neither entirely regular, nor are they entirely random. They are, quite simply, complex.

This book is about such collective animal behavior. It is a study of how interactions between animals produce group level patterns and why these interactions have evolved. It is aimed at two different types of reader: at the behavioural ecologist who is interested in how techniques for studying collective phenomena and complex systems are applied to animal groups and at the general scientist who would like to see solid examples of the application of techniques for the study of collective phenomena. I will start by arguing, for the sake of the bookshop browser who has at least read this far, why you should now buy this book.

1.1 Why collective animal behaviour?

The study of collective animal behavior gives a key example of how experiment can be combined with a theory of complex systems to better understand the world around us. As we have gained empirical knowledge of group behaviour across species there has been an increasing need for theoretical concepts, as well as mathematical and simulation models, that allow us to unify this knowledge into a general understanding. This book describes such concepts and tools, and shows how they can be applied.

Far from being a book solely about mathematical models, an equal focus is made on empirical examples. Each chapter describes a large number of different biological systems, and how one or two models can be applied in aiding our understanding of these different systems. This approach is taken in order to demonstrate the applicability of the theory to experiment and observation. Furthermore, by showing how similar models can be applied to very different systems I aim to demonstrate the logical relationships that can be drawn between very distinct biological systems. Modelling is a tool for understanding general properties of different systems.

Animal behaviour provides a wealth of interesting and accessible examples of collective phenomena and complex systems. Most people are familiar with ant trails; cockroach aggregations; fish schools; bird migrations; honey bee swarms; web construction by spiders; and locust marching, even if they have not observed them personally. In these systems there are two clearly defined levels of organisation which we aim to link together: the animal and the group. This clarity stands in contrast to many other collective phenomena, such as protein interactions or ecological webs, where it can be difficult to establish exactly on which level to observe a system. Thus animal behaviour provides much needed case studies of how complex systems theory can be put into practice.

What the study of animal behaviour might gain in terms of accessibility, it loses in terms of experimental precision. Animals are intrinsically more complicated than proteins or cells and it can be difficult to provide a clean description of the behaviour of individuals. Individual variation and difficulty in collecting large numbers of replicates means that we can seldom write down all-encompassing mathematical models for the actions and interactions of animals. The study of collective animal behaviour tests to the limit the supposed unifying nature of

mathematical modelling. Individual animals are inherently difficult to predict, but can we still make strong predictions about their collective behaviour? The answer is that in many situations we can. Applied correctly, mathematical models are able to make predictions about the group behaviour of even the most complex of animals, including humans.

The study of collective animal behaviour also allows us to better understand the various approaches to studying biology. The two main approaches are known as mechanistic and functional. The mechanistic approach looks at *how* animals interact to produce group level patterns. This approach concentrates on identifying communication mechanisms, such as visual and chemical signals, and tries to determine how these mechanisms are integrated to produce collective patterns. Traditionally, the study of collective phenomena and complex systems is more closely associated with such mechanistic explanations.

Functional explanations are based on arguments about *why* a behaviour evolved through natural selection: animals with behaviours that improve their chances of reproduction will increase in frequency in the population, and those with behaviours which are detrimental to survival will die out. Thus natural selection acts to produce behaviours consistent with the selfish interests of the individual or, more properly, the genes carried by the individual (Dawkins 1976; Dawkins 1982). Functional explanations take a central role in understanding why individuals co-operate to form collective patterns, why these patterns persist despite the conflicting interests of the individuals creating them, and why some collective patterns are inconsistent with the selfish individual or selfish gene.

Collective animal behaviour provides an excellent opportunity to study the link between function and mechanism. While these two forms of explanation are complementary, they are also interdependent. We cannot understand why co-operation evolves without knowing the mechanisms by which co-operative patterns are generated. Nor can we study mechanisms without considering the potential conflicts that can arise between individuals. This book seeks to clarify the relationship between mechanistic and functional explanations, by providing both sorts of explanation and showing how they fit together.

1.2 Mathematical modelling

The study of collective phenomena goes hand in hand with the use of mathematical and simulation modelling. The arguments made in this book often

rely on mathematics, and include applications of a wide range of mathematical models. What all of these models have in common is that they are logical arguments, taking us from one set of statements about the real world to another. Mathematics is a way of traveling logically from A to B , where A is a set of precisely stated rules or assumptions that are thought to characterize a system and B is a set of predictions logically arising from these properties. Mathematics is a language for making precise statements about the world and following their consequences to a logical conclusion (Feynman 1965 page 45).

In its idealized form, the mathematical modelling cycle proceeds as follows: (1) by formalizing some of our knowledge about a system in a mathematical model we generate a set of assumptions A ; (2) we use mathematical analysis or computer simulations to make a set of predictions B ; (3) we confirm or refute these predictions against our available knowledge and against the outcome of new experiments; (4) we return to step 1 and revise our assumptions in light of the outcome of the model and any new experiments. In theory, the cycle converges to a more and more accurate picture of reality.

Rather than being a strictly defined procedure, however, the modelling cycle is a convenient way of summarizing a whole range of mathematical and scientific activities. We are seldom in a position to consider or refine only a single set of assumptions or predictions. Instead, many different sets of assumptions can all have consequences consistent with the known relevant properties of a system. In such situations, modelling can suggest experiments which can discern which of several alternatives is the most accurate description of a system. It is often when we propose a model, make predictions, and find that these predictions are wrong that modelling provides its greatest insights into the real world.

Increased understanding of natural systems also arises from simply playing about with mathematical models. By looking at how different assumptions lead to different predictions and comparing these outcomes with what we know about a system, we can ensure that our own understanding of a system is logically consistent. Thus a great deal of worthwhile modelling remains at this stage of dividing our understanding of a system into the empirical observations which we take to be assumptions and the observations which we label as predictions. For example, we might ask the question: what is the minimal set of assumptions we can make that predicts all known system properties? This application of

Occam's razor serves to clarify and condense our understanding of a system in terms of the basic principles which underlie it.

It must be borne in mind that any particular mathematical model does not provide a unique way of thinking about the world. It is often the case that two completely different mathematical descriptions of a system are entirely compatible with each other. In particular, we can make two different sets of assumptions about a system, follow both of them to their logical conclusion, and produce two entirely different predictions. Provided the assumptions do not contradict one another, then neither do the predictions necessarily contradict each other. Remembering this apparently innocuous fact can resolve a lot of arguments about which of a variety of models of a system is 'best'.

1.3 From individuals to collectives

Models of collective animal behaviour are often based on assumptions about the behaviour at the level of individuals and then are used to make predictions about the patterns created at the level of the group. For example, we might describe how individual ants leave and follow chemical trails and predict the collective structure of their trail network. This division between the individual and the group leads us to expressions like 'emergence' and 'self-organisation' (Bonabeau et al. 1997; Camazine et al. 2001; Deneubourg & Goss 1989; Holland 1998; Kauffman 1993; Nicolis & Prigogine 1977). We make some relatively simple assumptions about individual behaviour from which emerge predictions about group behaviour. The group level pattern is said to self-organise because it was not encoded directly in the individual-level rules. There is thus a fundamental connection between mathematical modelling and self-organisation. Mathematical models are a way of extracting otherwise difficult-to-see connections between the interactions of individuals and the patterns created at a group level. The connection to mathematics removes any mystical meaning or ambiguity that might lurk within phrases such 'emergence' and 'self-organisation'. These phrases serve primarily to highlight the fact that simple interactions between individuals can produce sometimes surprising and empirically testable predictions about collective patterns.

There is a sense in which it is useful to talk about the principles arising from mathematical models of collective behaviour (Sumpter 2006). This is where the same model provides insights into many different and seemingly unrelated systems. Every chapter of this book provides an example where the same mathematical model has explanatory power across different biological systems.

For example, one model creates connections between firefly flashing and human applause (chapter 6) or another model connects ants foraging for food to cockroaches finding a shelter (chapters 3 and 4). This universal application of certain mathematical models is a remarkable observation. We can pick up the assumptions we have used to describe one system and apply them directly to produce predictions about a second system. Adding to this the fact that we can perturb our model and still produce the same predictions, or see changes in the predictions that reflect differences in the systems, would suggest fundamental laws encoded within these models.

These logical connections between systems prove extremely useful. In this book, I proceed on a case by case basis through different examples of collective animal behaviour. For each particular system, I classify how individuals interact with each other and make this the basis of mathematical models. When a similar mathematical model has been previously applied to another system, this helps us understand the behavioural algorithms and thus the system.

1.4 Functional and mechanistic approaches

Making assumptions about how individuals behave and predicting outcome at a group level is a mechanistic approach to the study of behaviour. We collect all the information about how individuals behave in response to their environment and to other individuals and incorporate this detail into a mathematical model that predicts the collective patterns generated by the group. In this way, we attempt to determine the mechanisms through which the collective outcome is formed.

The mechanistic approach to biology is often contrasted with the functional approach. In the functional approach we ask what the reproductive value or function is of a particular behavioural strategy. Mathematical models of function are usually based on assumptions about the costs and benefits, in terms of their impact on survival and reproduction, of a particular strategy. If we can identify the costs and the benefits associated with a strategy and compare them to those associated with an alternative strategy, we can predict how behaviour evolves through natural selection. We can also measure the extent to which animals are able to change their behavioural strategy in response to changing costs and benefits.

Functional questions are particularly interesting in the context of more than one individual, because what might be a benefit to one individual can be a cost to another. When natural selection acts to increase the frequency of a particular type of behavioural strategy in the population, it simultaneously changes the cost-benefit relationship for others in the population. As such, instead of acting to maximize some static function, natural selection acts to increase or decrease different types of strategy until an equilibrium is reached where no individual can evolve to do better by changing strategy (Dugatkin & Reeve 1998; Maynard-Smith 1982). This picture is complicated by interactions with relatives. If a genotype evolves which helps relatives then it can be selected for, not because it directly increases the survival value of the individual carrying the genotype, but because it increases the survival value of other individuals that also carry the genotype (Hamilton 1964).

Krebs & Davies (1993) provide a number of excellent examples of the difference between functional and mechanistic explanations. For example, female lions living in groups tend to come into oestrus at the same time. Krebs & Davies argue that there are functional benefits to the synchronisation of oestrus cycles. Mother lions suckle their young communally, so cubs born to a group of simultaneously lactating lions will have feeding opportunities even when their mother is out hunting. This benefit is further increased by the fact that the lions are relatives. Related lionesses pass on their genes not only through their own offspring, but also by raising the offspring of their sisters or cousins.

A mechanistic explanation looks at the process through which synchronization occurs. Lionesses' oestrus cycles are known to be coupled by the release of pheromones. When living alone, individual lionesses may have cycles with slightly different periods, but these cycles become entrained through the release of pheromones. Pheromone release by a lioness coming into oestrus will speed up or slow down those which are out of phase with her. Eventually, all lionesses will adopt the same cycle. This mechanistic explanation is well supported by mathematical models of synchronization (Strogatz 2003). These models predict that when group members behave periodically, but each with a different phase and possibly a different period, then, provided there is some means to communicate phase, the members can synchronise their cycles. The mathematical models give predictions about how factors such as individual variation and strength of coupling affect the degree of synchronisation.

Despite such clear examples, demarcating mechanistic and functional explanations can be a walk through a minefield. Let us for argument's sake consider a hypothetical study in which it was found that highly related females had synchronised oestrus cycles and less related females did not. It would be tempting to conclude from this study that unrelated females do not benefit from synchronization and thus do not synchronise their cycles. This is a functional prediction consistent with kin selected benefits. However, the mathematical results about synchrony state that if there is too much between individual variation in the period of cycles, as there may be between less related individuals, then synchronization becomes impossible (see chapter 6). The lack of synchrony may be due simply to a mechanistic 'failure' in unrelated groups. This alternative explanation could be partially resolved by, for example, measuring pheromone release or experimentally manipulating group composition to consist of unrelated lions with oestrus cycles with similar periods. Even after these tests there can remain problems in making functional interpretations. For example, did synchrony evolve because of kin selected benefits or do lion females remain in kin groups in order that they can synchronise their oestrus cycles?

Biologists are aware of these kinds of problems and I am confident that there would be some route through my rather simplified discussion of lion oestrus cycles. My point however is to re-emphasize that an increased knowledge of the mechanisms at work in any particular system can change our functional explanation of its behaviour. Mechanisms should not simply be considered as a way of obtaining parameters for the cost-benefit curves of functional models. Rather we should aim to form functional explanations that fully account for the underlying mechanisms. This point is particularly relevant where interactions produce highly non-linear patterns. Our intuition is not used to dealing with these outcomes and it becomes easy to miss important aspects of mechanisms that completely change our functional predictions. This book is full of examples of how mechanisms and functional explanations must be considered simultaneously if we are to come to a fuller understanding of group behaviour.

A more philosophical conflict sometimes arises between proponents of functional and mechanistic explanations. This conflict centres on the question of which type of explanation is more relevant to understanding biology. Factors influencing survival value are often described as 'ultimate', while 'proximate' is used to describe those which govern mechanisms. These labels can give the impression, and probably reflect the feelings of many behavioural ecologists, that functional questions have a greater importance than mechanisms (Krebs &

Davies 1997; West et al. 2007). This observation means that theories about self-organisation sometimes sit rather uncomfortably alongside the theory of evolution through natural selection. Throughout this book, I emphasize how the same mechanisms arise again and again in many different systems. Mathematical models formalize these logical connections between systems. For some scientists this gives these models and the principles which underlie them an equal, if not greater, importance than natural selection (Hoelzer et al. 2006; Kauffman 1993; Pepper & Hoelzer 2001; Wolfram 2002).

This book avoids such conflict, which usually obscures the real scientific questions. Whether verbal or mathematical, functional and mechanistic models are based on different assumptions and different predictions. More often than not these assumptions and predictions are consistent with each other. In other cases different decisions have been made about what is relevant or irrelevant in the construction of the models. As emphasized in section 1.2, it is perfectly consistent to have different models of a particular system, each of which makes a different link between assumptions and predictions. My main criterion for the choice between functional or mechanistic explanations in a particular instance is the strength of their explanatory and predictive power. Ideally, both types of explanation should be possible and we should be able to see the link between them.

1.5 Human society

One reason why animal groups are such a popular subject for scientific study is the importance of social interactions in our own everyday experience. Humans are inherently social animals, whose activities exhibit many of the elements of co-operation and conflict found in other animal societies. These social activities are extremely important to us: they determine our economic welfare; they produce a great deal of emotional turmoil, often providing the main reasons for whether we are happy or not; they determine how we are governed and how we structure our workplaces; and they even determine simple every day activities, such as how long we have to wait in queues.

Can some of the techniques used to study collective animal behaviour be applied to understanding human societies? The answer is a qualified 'yes'. In narrowly defined social situations, such as in pedestrian movement and spectator crowds, some of the techniques used to understand collective animal behaviour can be applied to humans. In wider situations, such as consumer decision-making and the 'evolution' of fads and fashions, there could also be

applications. Recent studies have looked at how our tendencies to buy particular items, find employment, and even commit crime change with the behaviours of those around us. Many of the underlying dynamics of these processes are similar to those seen in animal groups and this book seeks to highlight how these similarities arise.

1.6 Book structure

Before exhibiting any form of collective behaviour, animals must first come together. Chapter 2 looks at how and why animal groups form and the size distribution of these groups. This chapter sets the stage for more detailed investigations of the behaviour of animals once they have formed groups.

One advantage of living in a group is information transfer. Chapter 3 investigates signals that social animals have evolved to share information about the presence and location of food, as well as cues which some individuals use to parasitise the information possessed by others. Chapter 4 shows that similar principles of information transfer underlie collective decision-making. During migration, cockroaches, ants and bees all use similar rules to decide whether to move to a new home or shelter. These rules allow for consensus, whereby individuals 'agree' where to move to and are able to choose the best of several options.

Some spectacular examples of collective behaviour, which almost everyone has seen at some time, are moving animal groups. Chapter 5 looks at how simple models can capture the motion of fish schools, bird flocks and insect swarms. Chapter 6 then turns to synchronisation in time, looking at how fireflies co-ordinate their flashing and opera audiences co-ordinate their clapping.

Chapter 7 also concentrates on how simple rules can produce complex spatial patterns, but this time in the context of construction. This chapter looks at how termites and ants can build structures, such as nest mounds and trail networks, that are many orders of magnitude larger than they are. Chapter 8 looks at how interactions can lead to congestion and segregation. I discuss how economic and social systems often self-regulate to avoid these pitfalls, but sometimes also fail. Similar ideas can be applied to traffic congestion in ants and humans.

The repeated mantra that simple behaviour by individuals can produce complex patterns can sometimes be taken too far. Animals are not just simple individual units; each possesses a great deal of behavioural and physiological complexity. In chapter 9, I describe some methods for dealing with such intrinsic complexity and modelling complicated networks of interactions.

Although throughout this book I try to emphasise both mechanistic and functional approaches to studying collective animal behaviour, it is often the former which take a larger role in each of the chapters. Chapter 10 readdresses this balance by returning to the examples elsewhere in the book and discussing how they might have evolved through natural selection. The main aim of this chapter is to show how a thorough mechanistic understanding can also clarify our ultimate functional understanding of biological systems.

1.7 Working with models

Each chapter of this book has two or three boxes providing descriptions of key mathematical models. These models are most often simplified versions of those found elsewhere in the literature. The simplification and detailed description in the boxes has two complementary aims. Firstly, I hope that by simplifying the models I can illustrate better their central points. Secondly, I hope that the reader will be inspired to investigate and learn more about these models.

Understanding mathematical models requires more than just reading through their description. It is important to play around with them and find out how they work for different parameter values and react to small changes. For this reason, I have made the code for simulating most of the models presented in this book available online. The simulations run in Matlab and can be downloaded from

www.collective-behavior.com

The website also contains links to the homepages of many of the researchers mentioned in the book and links to key references.