# Preface

June 10, 2021

There has been a recent explosion of interest in the development of autonomous systems, and rightfully so. It is a field that promises to have an impact in all of our lives. The promise of autonomous systems ranges over the improvement of the workforce to acting as first explorers in planetary exploration. We look forward to living. in a world in which drones deliver our Amazon parcels within a few hours of placing an order and artificial agents that assist our medical professionals to provide high quality treatment and care in face of complex morbidities. All these wonderful applications have one thing in common, *they are inherently complex*. To require a robot to chart an entire planet or a single artificial agent to diagnose any ailment would be near impossible. So how can we design our agents to achieve all these complex tasks?

To answer this question, we must look to the only source of inspiration that we have - Nature. How do ants construct incredibly elaborate nests connected by an intricate network of tunnels? How do predators capture prey which are twice their size? Indeed, how do we humans design experiments which can tell us about the structure of atoms, and the origins of the Universe? The answer to all of this, of course, is that *we work together*.

The interaction of multiple entities has the power to produce what the individual has no hope of achieving. A few individuals working together form a group. Adding more individuals to this group forms a society. And, with yet more addition, a civilisation emerges. Increasing the number of individuals also increases complexity of operating together to perform a function. In fact, the more elements in a system, the greater is its potential complexity.

It is clear, then, that in order for our autonomous agents to achieve our lofty goals, we must design them so that they work together. But how to do this? The theory of multi-agent systems hopes to answer this very question, and does so from a number of perspectives.

One of the most prevalent amongst these perspectives is that of Game Theory, which mathematises the interactions between multiple agents (or players) and aims to understand what the resolution of this interaction will be. Game Theory acts as a playground for various thought experiments which, in turn, shed light on real world questions. Such questions, of course, are not limited to computational design, but can also help us understand our own world. Why, for instance, is altruistic behaviour so common in nature, despite the widely held belief that agents are inherently selfish? Which species will eventually die out through the process of evolution, and which will thrive?

So much for Game Theory, but how does learning fit into the picture? As we will see, the central thesis of traditional Game Theory is that agents are selfish, rational and, to some degree, well informed. These assumptions allow economists and computer scientists to make the assumption that the behaviour of these interactions is stationary; namely, the agents play their most rational choice given the actions of the other agents and then they stay there. As such, we can analyse the system at its equilibrium (which itself is no trivial task) and trust that this corresponds to empirical behaviour. Yet, as we see time and time again, organisms are not rational. One does not need to look far to see that they are not always well informed about the state of their environment, the actions of their fellow players or the consequences of their actions. These assumptions must, therefore be lifted, which lands us in the territory of *learning*. The view of multi-agent learning is that agents who are involved in a game do not always make the most rational decision. In fact, they often do not know what the most rational decision is. Instead they must adapt according to the behaviour of the other agents, and to an ever adapting environment.

This, fundamentally, leads to non-stationary behaviour, as we will see in these lectures. As such, we cannot analyse them as systems in equilibrium but must instead view them as adaptive and dynamic systems. In this series of lectures, we will be exploring the myriad intersections between the fields of Dynamical Systems and Game Theory. Specifically, we will try to understand the foundations of the dynamical aspects of learning.