

Chapter 1

Introduction

(1)

“Men are all alike.

IN WHAT WAY

They’re always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I’m depressed much of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

It’s true. I am unhappy.

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

I need some help, that much seems certain

WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

Perhaps I could learn to get along with my mother.

TELL ME MORE ABOUT YOUR FAMILY

My mother takes care of me.

WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU

My Father.” [Weizenbaum 1966]

(2)

“ LEFTLEFT LEFT”

The parts in capital letters of the two previous dialogues were both generated by computers exploiting Artificial Intelligence or AI. In the first example the computer program, called Eliza, was interacting with a human. The second example is a ‘translation’ of a dialogue between two simulated robots trying to help each other navigate a maze. This second example comes from an experiment reported later in this thesis. The initial impression we get from these two examples is that the first conversation seems to have been derived from a considerably more intelligent system than the latter. In this thesis we shall argue that this is not the case even though the range of words available to the second system is restricted to one or two words.

Eliza, acting as a therapist, relied on a very simple language parser and a few canned phrases as stock responses to keywords. “Tell me more about your family” was a response to the word ‘mother’ appearing in the previous human’s interaction. Beside the connection between the *words* of ‘mother’ and ‘family’, there is no sense in which Eliza has any knowledge of what a

mother is or anything else it talked about. The only possible interactions with the world for Eliza were through its conversations.

The rather more limited robots from the second example had a much more restricted conversational repertoire but we shall see that far more meaning was carried in the single word. The robots came from a population of similar creatures that had adaptable behaviour. Robots from this population learned to use their primitive motors and sensors to navigate around their maze environment. While they were initially endowed with an ability to create noise, they managed to forge this into a system that allowed them to convey meaning between themselves. Meaning based on what they learned to discriminate with their sensors and achieve through behaviours driven by their motors.

In the next chapter we consider some of the motivations of AI and argue for why *nouvelle*, bottom up AI, and particularly evolutionary robotics, has arisen to account for some of the failing of more traditional symbolic AI. We will then consider some shortcomings with this new approach that were adequately dealt with by traditional AI and look for a way to integrate the benefits of the two. An argument is presented for how this might be achieved by incorporating language acquisition as a fundamental goal into the framework of evolutionary robotics.

Chapter 3 lets us take a closer look at signalling and language. We consider what properties define this behaviour and help us to identify its presence, we also try to understand the way in which it can develop. Historical difficulties from both the fields of evolutionary biology and linguistics are addressed and some consideration is made for how we can go about this development. This leads us to consider the existing work on developing bottom up language behaviour already undertaken within artificial intelligence and artificial life literature. An analysis of the initial work in this discipline is used to offer insight into the subsequent development of sub-disciplines. These predominantly fall into two categories, motivation and mechanism.

In chapter 4 we explore the problems that arise on account of developing signals integrated with the moving target of developing behaviour. We present one of the main stumbling blocks to cooperative language development, which is the requirement for both receiver and transmitter strategies to be developed before benefits are accrued from signalling behaviour. Some ways are suggested in which these problems can be overcome through particular design features of the control architectures. The next three chapters report on experiments that explore these approaches.

Chapter 5 provides a detailed description of the experimental arena in which all of these studies are conducted. This also includes details of the agents involved and their learning mechanism and control architectures. Experiments are conducted to investigate how signalling can arise through imitation behaviour. The stability and potential for development of such signalling systems is considered.

In chapter 6 we look at a method that allows signal development to be enhanced for more general problem domains where there is no assistance from sources of imitation. Babbling used by human infants is used as motivation for creating ‘noisy’ agents. We explore how the amount of possible signals available to the agents affects their ability to acquire signalling behaviour. This is followed by measurements of how the rate of exploration of new signalling strategies affects the speed of development and the stability of such signalling behaviour.

With the last set of experiments, chapter 7 lets us present and test a method in which the development of signalling strategies can be enhanced by a radical design modification from previous agents. By making the agents use externalised signals for their own control, they are able to ‘share their thoughts’ with partners and for a large set of problems this greatly enhances the rate at which signalling can be developed.

In the final chapter the contributions of the thesis are summarised.